

# 000 001 002 003 004 005 SIMULATOR AND EXPERIENCE ENHANCED DIFFUSION 006 MODEL FOR COMPREHENSIVE ECG GENERATION 007 008 009

010 **Anonymous authors**  
011 Paper under double-blind review  
012  
013  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033  
034  
035  
036  
037  
038  
039  
040  
041  
042  
043  
044  
045  
046  
047  
048  
049  
050  
051  
052  
053

## ABSTRACT

Cardiovascular disease (CVD) is a leading cause of mortality worldwide. Electrocardiograms (ECGs) are the most widely used non-invasive tool for cardiac assessment, yet large, well-annotated ECG corpora are scarce due to cost, privacy, and workflow constraints. Generating ECGs can be beneficial for the mechanistic understanding of cardiac electrical activity, enable the construction of large, heterogeneous, and unbiased datasets, and facilitate privacy-preserving data sharing. Generating realistic ECG signals from clinical context is important yet underexplored. Recent work has leveraged diffusion models for text-to-ECG generation, but two challenges remain: (i) existing methods often overlook the physiological simulator knowledge of cardiac activity; and (ii) they ignore broader, experience-based clinical knowledge grounded in real-world practice. To address these gaps, we propose **SE-Diff**, a novel physiological simulator and experience enhanced diffusion model for comprehensive ECG generation. SE-Diff integrates a lightweight ordinary differential equation (ODE)-based ECG simulator into the diffusion process via a beat decoder and simulator-consistent constraints, injecting mechanistic priors that promote physiologically plausible waveforms. In parallel, we design an LLM-powered experience retrieval-augmented strategy to inject clinical knowledge, providing more guidance for ECG generation. Extensive experiments on real-world ECG datasets demonstrate that SE-Diff improves both signal fidelity and text-ECG semantic alignment over baselines, proving its superiority for text-to-ECG generation. We further show that the simulator-based and experience-based knowledge also benefit downstream ECG classification.

## 1 INTRODUCTION

Cardiovascular disease (CVD) remains a leading cause of global mortality and morbidity (Roth et al., 2020). In clinical workflows, the 12-lead electrocardiogram (ECG)—the standard setup using 10 electrodes to derive 12 voltage traces—is ubiquitous, non-invasive, and low-cost for screening, triage, and longitudinal monitoring (Kligfield et al., 2007). While machine learning (ML) has advanced ECG interpretation, progress is constrained by limited access to large, well-annotated corpora, stringent privacy considerations around sharing protected health information, and the expense of expert labeling (Johnson et al., 2023; Goldberger et al., 2000). ECG generation provides a principled way to mitigate these barriers by expanding training data, enabling controlled curation, and decoupling model development from directly identifiable records (Zanchi et al., 2025). Concurrently, denoising diffusion probabilistic models (DDPMs) and score-based methods have established strong fidelity and coverage across modalities (Ho et al., 2020; Song et al., 2021), motivating their adaptation to medical time series and, specifically, text-conditioned ECG generation (Lai et al., 2025a).

Despite these advances, there are still two gaps limit the practical adoption of text-to-ECG generation. (i) *Missing physiological simulator knowledge*. Most diffusion models for ECG learn morphology and timing purely from data, with minimal incorporation of known cardiac physiological dynamics. Decades of physiological modeling have produced compact ordinary differential equation (ODE) simulators that yield realistic P–QRS–T morphologies and heartrate variability under controllable parameters (McSharry et al., 2003; Malik, 1996). Yet these simulators are rarely integrated as priors or constraints during diffusion training, leaving a disconnect between statistical generation and mechanistic plausibility. (ii) *Under-use of experience-based knowledge at scale*. Prior text-to-ECG works often condition on narrow patient metadata, but do not leverage broader *experience*

knowledge—case-based regularities distributed across large electronic health record (EHR) corpora. Retrieval-augmented generation (RAG) offers a principled means to inject such non-parametric knowledge to generators (Lewis et al., 2020), including via lexical retrieval schemes, yet its potential for conditioning medical time-series generation remains underexplored.

To address these challenges, we introduce SE-Diff, a conditional latent-diffusion framework that synthesizes comprehensive ECG waveforms from natural-language clinical descriptions. SE-Diff couples a lightweight ODE-based ECG simulator to the denoising dynamics through a beat decoder—reconstructing a QRS-aligned single cycle from the latent code—and simulator-consistent spectral and rate constraints, thereby injecting mechanistic priors that steer generation toward physiologically plausible signals. In parallel, an LLM-powered retrieval pipeline identifies clinically similar patients from EHRs, retrieves their ECG diagnoses and measurements, and distills them into a concise, physiologically grounded description that is fused with available metadata for conditioning. In summary, our main contributions are as follows:

**1 Problem Identification.** We identify the problem of generating realistic 10s, 12-lead ECG waveforms directly from natural-language clinical descriptions. We propose SE-Diff, which can incorporate various patient metadata (age, sex, heart rate, rhythm/conduction) as soft clinical constraints to steer morphology toward clinically meaningful generation.

**2 Simulator-Informed Diffusion.** SE-Diff is the first to integrate a lightweight ODE-based ECG simulator with latent diffusion. We introduce a beat decoder that reconstructs a single-cycle beat from the latent representation, injecting simulator-consistent mechanistic priors that guide the denoising process toward physiologically plausible waveforms.

**3 Experience Retrieval-Augmented Conditioning.** We design an LLM-powered retrieval pipeline that identifies clinically similar patients based on EHR data and retrieves ECG diagnoses and measurements. The LLM generates a concise, physiologically grounded description, which is fused with available metadata to form the conditioning context.

**4 Experimental Validation.** Across real-world ECG datasets, SE-Diff surpasses baselines in both signal fidelity and text-ECG semantic alignment. Ablations quantify the contribution of simulator-based and experience-based knowledge conditioning. We further show that SE-Diff improves downstream ECG classification when used for augmentation.

## 2 PRELIMINARIES

### 2.1 DENOISING DIFFUSION PROBABILISTIC MODELS

Denoising Diffusion Probabilistic Models (DDPMs) (Sohl-Dickstein et al., 2015; Ho et al., 2020) define a fixed forward Markov noising process that maps a clean sample  $x_0 \sim q(x_0)$  to Gaussian noise over  $T$  steps, and a parametric reverse process that approximately inverts it. With variance schedule  $\{\beta_t\}_{t=1}^T \subset (0, 1)$ , set  $\alpha_t = 1 - \beta_t$  and  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ . The forward chain is

$$q(x_{1:T} \mid x_0) = \prod_{t=1}^T q(x_t \mid x_{t-1}), \quad (1)$$

where  $q(x_t \mid x_{t-1}) := \mathcal{N}(\sqrt{\bar{\alpha}_t} x_{t-1}, \beta_t \mathbf{I})$ . This implies the closed form  $x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon_t$ , with  $\epsilon_t \sim \mathcal{N}(0, \mathbf{I})$ . The exact reverse-time posterior  $q(x_{t-1} \mid x_t)$  is intractable, so DDPMs approximate it with a Gaussian transition  $p_\theta(x_{t-1} \mid x_t) := \mathcal{N}(\mu_\theta(x_t, t), \Sigma_\theta(x_t, t))$ , where a neural network predicts either the forward noise  $\epsilon$ , the clean signal  $x_0$ , or the velocity  $v$ . Under the common noise-prediction parameterization with  $\epsilon_\theta(x_t, t)$ , the mean is

$$\mu_\theta(x_t, t) = \frac{1}{\sqrt{\alpha_t}} \left( x_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right), \quad (2)$$

and we set  $\Sigma_\theta(x_t, t) \in \{\beta_t \mathbf{I}, \sigma_t^2 \mathbf{I}\}$ . Training maximizes a variational lower bound on  $\log p_\theta(x_0)$  (Sohl-Dickstein et al., 2015), which in practice reduces to the simple loss (Ho et al., 2020) with optional step-dependent weights  $w_t$ :

$$\mathcal{L}_{\text{simple}}(\theta) = \mathbb{E}_{t \sim \mathcal{U}\{1 \dots T\}, x_0 \sim q, \epsilon \sim \mathcal{N}} \left[ w_t \left\| \epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, t) \right\|_2^2 \right]. \quad (3)$$

108 2.2 THE ECG PHYSIOLOGICAL SIMULATOR  
109

110 In a resting heart, the ECG follows the P–QRS–T sequence. To reproduce this morphology, [McSharry et al. \(2003\)](#) proposed a three–ODE “ECG simulator” that generates realistic P–QRS–T  
111 waves while allowing control of heart-rate statistics and HRV spectrum ([Malik & Camm, 1990](#)).  
112 The model evolves a 3D state  $(x(t), y(t), z(t))$ :  $(x, y)$  traverse a unit-radius limit cycle whose angle  
113 encodes cardiac phase (one revolution per beat), and  $z(t)$  is the ECG voltage given by excursions  
114 about this cycle. The governing ODEs are:  
115

$$116 \quad \frac{dx}{dt} = \alpha(x, y) x - \omega y, \quad \frac{dy}{dt} = \alpha(x, y) y + \omega x. \quad (4)$$

$$119 \quad \frac{dz}{dt} = - \sum_{\beta \in \{P, Q, R, S, T\}} a_{\beta} \Delta \theta_{\beta}(x, y) \exp\left(-\frac{\Delta \theta_{\beta}(x, y)^2}{2 b_{\beta}^2}\right) - [z - z_0(t)]. \quad (5)$$

120 where  $\alpha(x, y) = 1 - \sqrt{x^2 + y^2}$  drives  $(x, y)$  toward the unit limit cycle,  $\theta(x, y) = \text{atan2}(y, x) \in$   
121  $[-\pi, \pi]$  is the phase, and  $\Delta \theta_{\beta}(x, y) = (\theta(x, y) - \theta_{\beta}) \bmod 2\pi$  is the phase offset to landmark  $\beta \in \mathcal{B}$   
122 with  $\mathcal{B} = \{P, Q, R, S, T\}$ . The parameter  $\omega$  controls angular velocity (thus average heart rate), and  
123  $z_0(t)$  is a slow baseline (e.g., respiratory wander modeled as  $z_0(t) = A \sin(2\pi f_{\text{resp}} t)$  with small  
124 amplitude  $A$  ([Sörnmo & Laguna, 2005](#))). All morphology parameters are collected as  
125

$$127 \quad \eta = \{\theta_P, \theta_Q, \theta_R, \theta_S, \theta_T, a_P, a_Q, a_R, a_S, a_T, b_P, b_Q, b_R, b_S, b_T\}, \quad (6)$$

128 where these parameters—phase landmarks  $\theta_{\beta}$ , amplitude coefficients  $a_{\beta}$ , and width coefficients  $b_{\beta}$   
129 for each  $\beta \in \{P, Q, R, S, T\}$ —govern the ECG morphology. When the  $(x, y)$  state passes the phase  
130  $\theta_{\beta}$ , the Gaussian-shaped term  $a_{\beta} \Delta \theta_{\beta} \exp(-\Delta \theta_{\beta}^2/(2b_{\beta}^2))$  in 5 transiently perturbs  $z$  away from  
131 baseline, producing the corresponding P/QRS/T deflection. The sign of  $a_{\beta}$  sets polarity (upward for  
132  $a_{\beta} > 0$ , downward for  $a_{\beta} < 0$ );  $|a_{\beta}|$  controls peak amplitude; and  $b_{\beta}$  sets the temporal spread  
133 (wave duration). The restoring term  $-[z - z_0(t)]$  then pulls the signal back toward baseline. Unless  
134 otherwise specified, we adopt the parameter values recommended by [McSharry et al. \(2003\)](#).  
135

136 **The Euler Method.** To simulate the synthetic ECG  $z(t)$ , we numerically solve the ODE system  
137 with a fixed-step explicit Euler method (the first-order Runge–Kutta scheme) ([Butcher & Butcher,  
138 1987; Süli & Mayers, 2003](#)). We choose the step size  $\Delta t = 1/f_s$  to match the desired sampling  
139 frequency (e.g.,  $f_s = 500$  Hz). Using the finite-difference approximation ([Milne-Thomson, 2000](#)):  
140

$$140 \quad \frac{du}{dt}(t) \approx \frac{u(t + \Delta t) - u(t)}{\Delta t}, \quad (7)$$

142 which leads to the update rule  $u(t + \Delta t) = u(t) + v(t) \Delta t$ , for an ODE of the form  $du/dt = v(t)$ .  
143 Starting from initial conditions  $(x_0, y_0, z_0)$ , we iterate this update for each time step. At the  $\ell$ -th  
144 step (time  $t_{\ell} = \ell \Delta t$ ), let  $v_{\ell} = (f_x(x_{\ell}, y_{\ell}; \eta), f_y(x_{\ell}, y_{\ell}; \eta), f_z(x_{\ell}, y_{\ell}, z_{\ell}, t_{\ell}; \eta))$  denote the  
145 right-hand side of Equation 4 and 5. The state is then advanced as:  
146

$$147 \quad x_{\ell+1} = x_{\ell} + f_x(x_{\ell}, y_{\ell}; \eta) \Delta t, \quad (8)$$

$$148 \quad y_{\ell+1} = y_{\ell} + f_y(x_{\ell}, y_{\ell}; \eta) \Delta t, \quad (9)$$

$$149 \quad z_{\ell+1} = z_{\ell} + f_z(x_{\ell}, y_{\ell}, z_{\ell}, t_{\ell}; \eta) \Delta t, \quad (10)$$

151 and this process is repeated for  $\ell = 0, 1, 2, \dots$  up to the desired number of samples  $L$ . In other  
152 words, each iteration uses the derivatives  $f_x, f_y, f_z$  at the current state to step the solution forward  
153 by  $\Delta t$ . This simple explicit scheme is computationally efficient and sufficient for our purposes,  
154 though higher-order integration methods could be used for greater accuracy if needed.  
155

156 3 METHOD: SE-DIFF  
157

158 We present SE-Diff, a conditional latent-diffusion framework that synthesizes 10s, 12-lead ECGs  
159 from clinical text. Diffusion operates in the VAE latent space (Sec. 2.1). To make physiology-  
160 aware supervision tractable, we attach a lightweight *Beat Decoder* that predicts a single QRS-aligned  
161 cardiac cycle from the latent; its output drives simulator-informed regularizers derived from the ECG  
physiology model in Sec. 2.2. To strengthen conditioning, SE-Diff also incorporates experience

knowledge retrieved based on EHRs (Sec. 3.4). In inference, we sample in latent space and decode with the full VAE.

**Problem Formulation.** Each ECG record is a multivariate sequence  $\mathbf{x} \in \mathbb{R}^{12 \times L}$  representing a 10s, 12-lead waveform sampled at  $f_s$ . Our goal is to learn a conditional generator  $p(\mathbf{x} | c; \phi, \vartheta, \theta)$  that uses  $c$  throughout denoising to produce physiologically plausible ECGs. The conditioning is  $c = (t, m, r)$ , comprising original diagnoses  $t$ , basic metadata  $m$  (age, sex, optionally heart rate), and retrieve-augmented report  $r$ . Concretely, we first train a VAE ( $E_\phi, D_\theta$ ) together with a *Beat Decoder*  $D_\psi^{\text{beat}}$ ; the encoder maps a full recording to a latent sequence  $z_0 = E_\phi(\mathbf{x}) \in \mathbb{R}^{d \times T}$ , where  $T = L/S$  and  $S$  is the VAE temporal stride, and  $D_\psi^{\text{beat}}$  maps  $z_0$  to a single-cycle prediction  $h \in \mathbb{R}^{12 \times L_c}$ . We then freeze  $E_\phi$ ,  $D_\theta$ , and  $D_\psi^{\text{beat}}$  and train a DDPM in latent space using a U-Net denoiser  $\epsilon_\vartheta(z_t, t, c)$  with cross-attention to  $c$ . During diffusion training, simulator-guided penalties (Sec. 3.3) are applied to the Beat Decoder output  $D_\psi^{\text{beat}}(z_0)$ , while experience–knowledge features augment the text pathway (Sec. 3.4). At test time, we run the reverse process to obtain  $\hat{z}_0$  and decode  $\hat{\mathbf{x}} = D_\theta(\hat{z}_0)$  (Sec. 3.5).

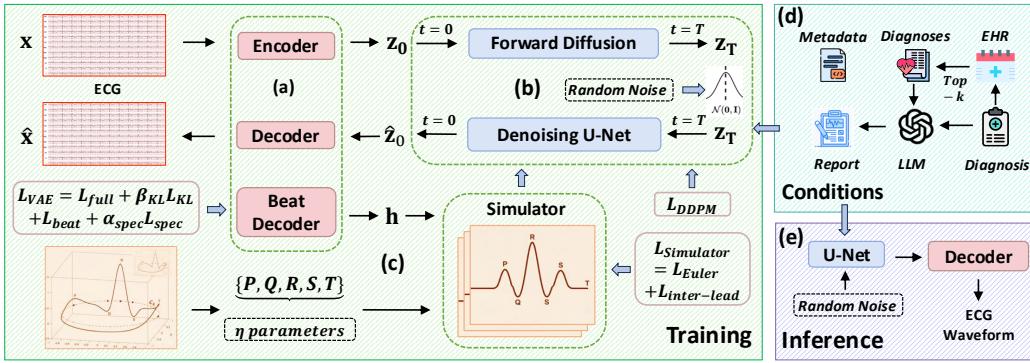


Figure 1: **Overview Framework of SE-Diff.** (a) *Variational Autoencoder*: encoder–decoder with a lightweight beat decoder for a single QRS-aligned cycle. (b) *Conditional latent diffusion*: U-Net denoiser with cross-attention to text, metadata, and retrieved report. (c) *Simulator-informed diffusion*: Euler and inter-lead constraints on the beat decoder output. (d) *Experience retrieval-augmented Conditioning*: tri-view EHR similarity with LLM distillation into a concise report. (e) *Inference*: reverse diffusion and decoding to a 10s, 12-lead ECG.

### 3.1 VARIATIONAL AUTOENCODER

We learn a latent representation for 12-lead ECGs with a variational autoencoder (VAE). Given a full recording  $\mathbf{x} \in \mathbb{R}^{12 \times L}$ , the encoder  $E_\phi$  parameterizes a diagonal Gaussian posterior

$$q_\phi(z | \mathbf{x}) = \mathcal{N}(z; \mu_\phi(\mathbf{x}), \text{diag}(\sigma_\phi^2(\mathbf{x}))), \quad z_0 = \mu_\phi(\mathbf{x}) + \sigma_\phi(\mathbf{x}) \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (11)$$

where  $z_0 \in \mathbb{R}^{d \times T}$  with  $T = L/S$  and  $S$  the VAE temporal stride. The decoder  $D_\theta$  reconstructs the signal  $\hat{\mathbf{x}} = D_\theta(z_0) \in \mathbb{R}^{12 \times L}$ . To expose morphology at the beat scale, we attach a lightweight Beat Decoder:  $D_\psi^{\text{beat}} : \mathbb{R}^{d \times T} \rightarrow \mathbb{R}^{12 \times L_c}$  to get the single cycle signal:  $h = D_\psi^{\text{beat}}(z_0)$ .

**Training.** Let  $r_0$  denote the first R-peak index at sampling rate  $f_s$  (Golany et al., 2020b); define  $\mathcal{C}(\mathbf{x}) = \mathbf{x}[:, r_0 - 0.2 f_s : r_0 + 0.4 f_s]$  and  $L_c = 0.2 f_s + 0.4 f_s$ . We train the encoder, decoder and Beat Decoder with length-normalized mean-squared errors (MSE) and a single KL term:

$$\mathcal{L}_{\text{full}} = \frac{1}{12L} \|\mathbf{x} - D_\theta(E_\phi(\mathbf{x}))\|_F^2, \quad \mathcal{L}_{\text{KL}} = \text{KL}(q_\phi(z | \mathbf{x}) \| \mathcal{N}(0, \mathbf{I})), \quad (12)$$

$$\mathcal{L}_{\text{beat}} = \frac{1}{12L_c} \|\mathcal{C}(\mathbf{x}) - D_\psi^{\text{beat}}(E_\phi(\mathbf{x}))\|_F^2. \quad (13)$$

The Beat Decoder’s single beat should also reflect the statistics of all beats in the 10s window. Rather than tiling  $h$  to length  $L$ , we detect all R-peaks within the window,  $\{r_j\}_{j=1}^J$  (with  $J$  determined by the number of detected beats), and extract per-beat crops of identical length  $L_c$ :

216  $\mathcal{C}_j(\mathbf{x}) = \mathbf{x}[:, r_j - 0.2 f_s : r_j + 0.4 f_s] \in \mathbb{R}^{12 \times L_c}$ ,  $j = 1, \dots, J$ . Let the Beat Decoder output  
 217 be  $h = D_\psi^{\text{beat}}(z_0) \in \mathbb{R}^{12 \times L_c}$ . For each lead  $\ell$  and each cycle  $j$ , we remove the mean and compute  
 218 the one-sided log-magnitude spectrum (real FFT of length  $L_c$ ) up to  $f_{\max}$ :  
 219

$$220 \quad S_\ell^{(j)}[k] = \log(\varepsilon + |\text{rFFT}(\mathcal{C}_j(\mathbf{x})_\ell - \overline{\mathcal{C}_j(\mathbf{x})_\ell})[k]|), \quad \hat{S}_\ell[k] = \log(\varepsilon + |\text{rFFT}(h_\ell - \bar{h}_\ell)[k]|), \quad (14)$$

221 with  $\varepsilon > 0$  small and frequencies  $f_k = \frac{k}{L_c} f_s$ . The spectral loss averages the (optionally  
 222 band-weighted) squared discrepancy between the Beat Decoder’s spectrum and the spectrum of  
 223 each observed cycle:  
 224

$$225 \quad \mathcal{L}_{\text{spec}} = \frac{1}{12JK} \sum_{\ell=1}^{12} \sum_{j=1}^J \sum_{k: f_k \leq f_{\max}} w(f_k) (\hat{S}_\ell[k] - S_\ell^{(j)}[k])^2, \quad (15)$$

228 where  $K = |\{k : f_k \leq f_{\max}\}|$  and  $w(f)$  can emphasize physiologically salient bands (e.g., higher  
 229 weights on 0.5–3 Hz for heart rate). We jointly optimize the encoder, the full decoder, and the beat  
 230 decoder with:  
 231

$$232 \quad \mathcal{L}_{\text{VAE}} = \mathcal{L}_{\text{full}} + \beta_{\text{KL}} \mathcal{L}_{\text{KL}} + \mathcal{L}_{\text{beat}} + \alpha_{\text{spec}} \mathcal{L}_{\text{spec}}, \quad (16)$$

233 where  $\alpha_{\text{spec}} > 0$ . Length normalization makes  $\mathcal{L}_{\text{full}}$  and  $\mathcal{L}_{\text{beat}}$  commensurate;  $\mathcal{L}_{\text{spec}}$  incorporates  
 234 full-window frequency statistics into the single-cycle prediction.  
 235

### 3.2 CONDITIONAL LATENT DIFFUSION

237 We train the diffusion in the VAE latent space. Given a latent sequence  $z_0 = E_\phi(\mathbf{x}) \in \mathbb{R}^{d \times T}$ , the  
 238 forward process follows Sec. 2.1 with  $x_t \mapsto z_t$ :  $z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ , and  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ . We  
 239 train a conditional denoiser  $\epsilon_\vartheta(z_t, t, c)$  with the standard objective  
 240

$$241 \quad \mathcal{L}_{\text{DDPM}} = \mathbb{E}_{t, z_0, \epsilon} \left[ \left\| \epsilon - \epsilon_\vartheta(z_t, t, c) \right\|_2^2 \right]. \quad (17)$$

243 The denoiser is a 1D U-Net (Ronneberger et al., 2015) that treats the latent as a sequence  $z_t \in \mathbb{R}^{d \times T}$   
 244 (channels  $d$ , length  $T$ ). Conditioning enters via cross-attention to a context representation  $C$  built  
 245 from  $c = (t, m, r)$  (Rombach et al., 2022; Vaswani et al., 2017); a final  $1 \times 1$  convolution maps fea-  
 246 tures to  $\epsilon_\vartheta$ . Timestep and context embeddings modulate intermediate features through FiLM-style  
 247 affine transformations (Perez et al., 2018). At sampling, we use standard DDPM transitions with  
 248 classifier-free guidance (Ho & Salimans, 2022) and common improvements such as cosine sched-  
 249 ules and optional learned variances (Nichol & Dhariwal, 2021); further architectural and training  
 250 details are provided in the appendix.  
 251

### 3.3 SIMULATOR-INFORMED DIFFUSION

254 We estimate class-specific parameters  $\eta_{\text{class}}$  offline by fitting the simulator to representative real  
 255 beats using an explicit Euler integrator, together with lightweight stabilizers and morphology pri-  
 256 ors that improve convergence and preserve physiological plausibility (details in Appendix C). For  
 257 each training sample, we obtain a single-cycle waveform  $h = D_\psi^{\text{beat}}(z_0) \in \mathbb{R}^{12 \times L_c}$  from the Beat  
 258 Decoder (Sec. 3.1). This beat is used to enforce mechanistic plausibility via an ODE-based ECG  
 259 simulator (Sec. 2.2). The simulator’s morphology parameters  $\eta = \{\theta_\beta, a_\beta, b_\beta\}_{\beta \in \{P, Q, R, S, T\}}$  enter  
 260 the right-hand side  $f_z(\cdot; \eta)$  of Eq. 5, which defines  $f_z$ . During diffusion training, the simulator  
 261 provides two complementary regularizers:  
 262

**Simulator-guided Euler Loss.** Given the single-cycle waveform  $h \in \mathbb{R}^{12 \times L_c}$ , we integrate the  
 263 simulator with parameters  $\eta$  and fixed initials  $(x_0, y_0)$  to obtain a reference trajectory  $(x_t, y_t)$  at  
 264  $t = \ell \Delta t$ . We penalize per-lead deviations from the ODE and the simulator-guided Euler loss is:  
 265

$$266 \quad \mathcal{L}_{\text{Euler}} = \frac{1}{12(L_c - 1)} \sum_{\text{lead}} \sum_{\ell=1}^{L_c-1} \left( \frac{h_{\ell+1} - h_\ell}{\Delta t} - f_z(x_\ell, y_\ell, h_\ell, t_\ell; \eta) \right)^2. \quad (18)$$

268 **Inter-lead dependency constraint.** Realistic 12-lead synthesis requires not only accurate per-lead  
 269 morphology but also correct physiological interdependencies among leads. We therefore enforce the  
 270

270 classical frontal-plane identities implied by the standard ECG configuration (Einthoven’s triangle  
 271 and Goldberger’s central terminal), constraining the generated limb and augmented leads to remain  
 272 mutually consistent:

$$\begin{aligned} I &= II - III, \quad aVR = -\frac{1}{2}(I + II), \quad aVL = \frac{1}{2}(I - III), \\ II &= I + III, \quad aVF = \frac{1}{2}(II + III), \quad III = II - I. \end{aligned} \quad (19)$$

273 Let  $\mathcal{L}_{\text{frontal}} = \{I, II, III, aVR, aVL, aVF\}$  be the frontal-plane leads (a subset of the 12 leads).  
 274 For any identity of the form  $y = \beta p + \gamma q$  with  $y, p, q \in \mathcal{L}_{\text{frontal}}$ , we treat  $y$  as the *child* lead and  
 275  $p, q$  as its *parent* leads. Denote by  $h_\ell^L$  the sample at index  $\ell$  of lead  $L$  from the predicted 12-lead  
 276 beat  $h \in \mathbb{R}^{12 \times L_c}$ . We obtain parent simulator states  $(x_\ell^p, y_\ell^p)$  and  $(x_\ell^q, y_\ell^q)$  by integrating the sim-  
 277 ular with class-specific parameters  $\eta_p, \eta_q$ . Defining  $\mathcal{C} = \{(I, II, III, 1, -1), (II, I, III, 1, 1),$   
 278  $(III, II, I, 1, -1), (aVR, I, II, -\frac{1}{2}, -\frac{1}{2}), (aVL, I, III, \frac{1}{2}, -\frac{1}{2}), (aVF, II, III, \frac{1}{2}, \frac{1}{2})\}$  and the  
 279 loss aggregates the six constraints over time:

$$\mathcal{L}_{\text{inter-lead}} = \sum_{(y, p, q, \beta, \gamma) \in \mathcal{C}} \sum_{\ell=1}^{L_c-1} \left( \frac{h_{\ell+1}^y - h_\ell^y}{\Delta t} - \beta f_z(x_\ell^p, y_\ell^p, h_\ell^p, t_\ell; \eta_p) - \gamma f_z(x_\ell^q, y_\ell^q, h_\ell^q, t_\ell; \eta_q) \right)^2. \quad (20)$$

280 Here  $y, p, q$  are specific leads (elements of the 12-lead set), and  $\mathcal{C}$  enumerates each frontal-plane  
 281 identity as a tuple  $(y, p, q, \beta, \gamma)$ . This construction directly matches the child’s discrete derivative to  
 282 the corresponding linear combination of the parents’ simulator derivatives, enforcing physiologically  
 283 grounded inter-lead consistency.

### 284 3.4 EXPERIENCE RETRIEVAL–AUGMENTED CONDITIONING

285 We augment text conditioning with clinical experience retrieved from electronic health records  
 286 (EHR). Specifically, we link MIMIC-IV-ECG (Gow et al., 2023) to MIMIC-IV-CLINICAL (Johnson et al., 2023), build a compact tri-view profile (diagnoses, medications, procedures) (Ou et al., 2025), and retrieve the top- $k$  clinically similar admissions. For an index admission  $u$ , let  $E_u^{\text{Diag}}$ ,  
 287  $E_u^{\text{Med}}$ , and  $E_u^{\text{Proc}}$  denote the sets of diagnosis, medication, and procedure codes, respectively.  
 288 Given another admission  $u'$ , we compute set similarities using the Jaccard index  $J(A, B)$  for  
 289  $X \in \{\text{Diag, Med, Proc}\}$ :

$$\tau_X(u, u') = J(E_u^X, E_{u'}^X). \quad (21)$$

300 These similarities are aggregated with nonnegative weights  $\lambda_1, \lambda_2, \lambda_3$  to yield a single similarity:

$$\tau(u, u') = \lambda_1 \tau_{\text{Diag}}(u, u') + \lambda_2 \tau_{\text{Med}}(u, u') + \lambda_3 \tau_{\text{Proc}}(u, u'). \quad (22)$$

301 We retrieve the top- $k$  most similar admissions according to the tri-view similarity over diagnoses,  
 302 medications, and procedures, and then pass their diagnostic profiles together with  $(t, m)$  to the LLM  
 303 using the prompt shown in Fig. 5 in the Appendix to obtain a concise, physiologically grounded  
 304 report  $r$ , where  $t$  denotes the original diagnoses and  $m$  encodes basic metadata (age, sex, option-  
 305 ally heart rate). Finally, the conditioning input is  $c = (t, m, r)$ , which conditions the denoiser via  
 306 cross-attention.

### 307 3.5 TRAINING AND INFERENCE

308 **Training objective.** We combine the latent-space diffusion loss with simulator-based regularizers:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{DDPM}} + \lambda \mathcal{L}_{\text{Euler}} + \gamma \mathcal{L}_{\text{inter-lead}}, \quad \lambda, \gamma > 0. \quad (23)$$

309 We first train the VAE using Eq. 16 and then freeze  $E_\phi$ ,  $D_\theta$ , and  $D_\psi^{\text{beat}}$ . During diffusion training,  
 310 we optimize only the denoiser  $\epsilon_\vartheta$ ; the Beat Decoder appears only through these regularizers—we  
 311 use  $D_\psi^{\text{beat}}$  to produce  $h = D_\psi^{\text{beat}}(z_0)$  for  $\mathcal{L}_{\text{Euler}}$  and  $\mathcal{L}_{\text{inter-lead}}$ . All simulator-driven terms are  
 312 training-only and do not modify the reverse process.

313 **Inference.** Given conditioning  $c$ , we draw  $z_T \sim \mathcal{N}(0, \mathbf{I})$  and apply the learned reverse diffusion  
 314 from  $t = T$  to 1 with the standard DDPM parameterization (variance schedule  $\{\beta_t\}$ ,  $\alpha_t = 1 - \beta_t$ ,  
 315  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ ):

$$\hat{z}_0(z_t, t, c) = \frac{z_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\vartheta(z_t, t, c)}{\sqrt{\bar{\alpha}_t}}, \quad (24)$$

$$\mu_\theta(z_t, t, c) = \frac{1}{\sqrt{\alpha_t}} \left( z_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(z_t, t, c) \right). \quad (25)$$

We set  $\tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t$  and sample  $z_{t-1} = \mu_\theta(z_t, t, c) + \sqrt{\tilde{\beta}_t} \xi_t$  with  $\xi_t \sim \mathcal{N}(0, \mathbf{I})$ . After the final step, we decode to the signal domain,  $\hat{\mathbf{x}} = D_\theta(\hat{z}_0)$ , optionally using classifier-free guidance during sampling.

## 4 EXPERIMENTS

**Dataset and Preprocessing.** We train on MIMIC-IV-ECG (Gow et al., 2023; Johnson et al., 2023), which contains 800,035 de-identified 12-lead, 10 s ECGs sampled at 500 Hz. Heart rate (HR) is taken from metadata when available; otherwise it is re-estimated via QRS detection (WFDB XQRS). Waveforms are encoded by a VAE into  $4 \times 128$  latents that serve as inputs to the diffusion model. We use the MIMIC-IV-Clinical (Johnson et al., 2023) to obtain each patient’s EHR for experience knowledge conditioning. For external validation, we additionally experiment on the PTB-XL dataset (Wagner et al., 2020), which consists of 12-lead clinical ECGs with standardized diagnostic labels.

**Baselines.** We compare against four strong ECG generation baselines: (i) *DiffuSETS* (Lai et al., 2025a), to our knowledge the only prior method that generates *12-lead, 10 s* ECGs from clinical text; (ii) a GAN-based model (WGAN) originally proposed for arrhythmia classification using cGAN-augmented ECG signals (Adib et al., 2022), which we adapt to synthesize 10 s, 12-lead ECGs under our setting; (iii) SSSD, a diffusion-based conditional ECG generator built on structured state space models (López Alcaraz & Strothoff, 2023); and (iv) BeatDiff, an ECG beat diffusion model designed for morphology-aware reconstruction from indirect signals (Bedin & Coauthors, 2024). To quantify the contribution of each component of SE-Diff, we report ablations trained under identical schedules and seeds: (i) *SE-Diff w/o Sim* (removing the Euler consistency term  $\mathcal{L}_{\text{Euler}}$ ); (ii) *SE-Diff w/o InterLead* (dropping  $\mathcal{L}_{\text{inter-lead}}$ ); (iii) *SE-Diff w/o Exp* (disabling EHR retrieval and LLM distillation so conditioning uses only text+metadata).

### 4.1 ECG GENERATION RESULTS

We evaluate SE-Diff along four clinically aligned levels: *signal-level stability, feature-level physiology, diagnostic/semantic alignment, and beat-level morphology and interval fidelity*. At each level, we define the metrics and report aggregate results on both MIMIC-IV-ECG and PTB-XL.

**Signal-level Stability.** Given matched real and generated ECGs ( $\mathbf{x}, \hat{\mathbf{x}}$ ) under the same condition  $c$ , we compute per-lead mean absolute error (MAE), normalized root mean squared error (NRMSE), and Pearson correlation  $r$  to assess waveform fidelity and temporal consistency.

**Feature-level Physiology.** To evaluate preservation of basic physiology, we compare heart rate (HR) estimated from  $\hat{\mathbf{x}}$  and  $\mathbf{x}$  via the absolute error  $\text{MAE}_{\text{HR}}$ .

**Diagnostic Alignment.** We adopt a CLIP-style evaluation for ECG–text pairs: an ECG encoder  $f_{\text{ecg}}(\cdot)$  and a text encoder  $f_{\text{text}}(\cdot)$  produce  $\ell_2$ -normalized embeddings; cosine similarity quantifies alignment,  $s(\mathbf{x}, \text{text}) = \langle f_{\text{ecg}}(\mathbf{x}), f_{\text{text}}(\text{text}) \rangle$ . To control encoder bias, we report the relative CLIP score and the relative Fréchet distance:

$$\text{rCLIP} = \frac{s(\hat{\mathbf{x}}, \text{text})}{s(\mathbf{x}, \text{text})} \quad \text{rFID} = \frac{\text{FID}(\hat{\mathcal{X}}, \mathcal{X}_r)}{\text{FID}(\mathcal{X}_r^{(1)}, \mathcal{X}_r^{(2)})}. \quad (26)$$

Distributional coverage/quality is measured with the Fréchet distance in the ECG embedding space,  $\text{FID} = \|\mu_r - \mu_g\|_2^2 + \text{Tr}(\Sigma_r + \Sigma_g - 2(\Sigma_r \Sigma_g)^{1/2})$ , where  $(\mu_r, \Sigma_r)$  and  $(\mu_g, \Sigma_g)$  denote the mean and covariance of real and generated ECG embeddings,  $\hat{\mathcal{X}}$  denotes generated samples, and  $\mathcal{X}_r^{(1)}, \mathcal{X}_r^{(2)}$  are disjoint splits of the real set.

**Beat-level Morphology and Interval Fidelity.** To assess beat-level fidelity beyond global signal errors and heart rate, we introduce a set of quantitative morphology and interval features. For each ECG, we automatically extract PR interval, QRS duration (QRSd), QT interval, heart-rate–corrected

378 Table 1: ECG generation performance on MIMIC-IV-ECG and PTB-XL datasets.  
379

380 Model	381 MAE $\downarrow$	381 NRMSE $\downarrow$	381 MAE <sub>HR</sub> $\downarrow$	381 rCLIP Score $\uparrow$	381 rFID Score $\uparrow$
<b>MIMIC-IV-ECG (internal)</b>					
382 SSDM	383 0.4337 $\pm$ 0.0300	383 0.2027 $\pm$ 0.0441	383 27.37 $\pm$ 14.84	383 0.7213 $\pm$ 0.0402	383 0.9096 $\pm$ 0.0398
383 WGAN	384 0.1896 $\pm$ 0.0605	384 0.1301 $\pm$ 0.0316	384 31.54 $\pm$ 2.15	384 0.5688 $\pm$ 0.0347	384 0.5497 $\pm$ 0.0192
384 BeatDiff	385 0.7464 $\pm$ 0.0070	385 0.4756 $\pm$ 0.0117	385 27.74 $\pm$ 1.49	385 0.5167 $\pm$ 0.0180	385 0.8612 $\pm$ 0.0039
385 DiffuSETS	386 0.1092 $\pm$ 0.0022	386 0.0851 $\pm$ 0.0012	386 13.29 $\pm$ 1.13	386 0.9309 $\pm$ 0.0036	386 0.9209 $\pm$ 0.0038
<b>PTB-XL (external)</b>					
390 SSDM	391 0.6103 $\pm$ 0.0204	391 0.3818 $\pm$ 0.0670	391 15.22 $\pm$ 11.51	391 0.8618 $\pm$ 0.0599	391 0.7168 $\pm$ 0.0355
391 WGAN	392 0.2458 $\pm$ 0.0653	392 0.1197 $\pm$ 0.0313	392 13.82 $\pm$ 18.69	392 0.5880 $\pm$ 0.0000	392 0.5377 $\pm$ 0.0232
392 BeatDiff	393 0.9888 $\pm$ 0.0059	393 0.4731 $\pm$ 0.0104	393 13.86 $\pm$ 0.78	393 0.8799 $\pm$ 0.0022	393 0.8503 $\pm$ 0.0035
393 DiffuSETS	394 0.1281 $\pm$ 0.0030	394 0.0797 $\pm$ 0.0011	394 17.88 $\pm$ 0.72	394 0.8690 $\pm$ 0.0011	394 0.8456 $\pm$ 0.0035
<b>SE-Diff (ours)</b>					
394 <b>SE-Diff (ours)</b>	395 <b>0.1076<math>\pm</math>0.0033</b>	395 <b>0.0630<math>\pm</math>0.0006</b>	395 <b>8.24<math>\pm</math>0.43</b>	395 <b>0.8901<math>\pm</math>0.0060</b>	395 <b>0.8583<math>\pm</math>0.0056</b>
395 w/o Sim	396 0.1138 $\pm$ 0.0032	396 0.0680 $\pm$ 0.0007	396 14.72 $\pm$ 0.90	396 0.8896 $\pm$ 0.0010	396 0.8004 $\pm$ 0.0061
396 w/o InterLead	397 0.1084 $\pm$ 0.0034	397 0.0640 $\pm$ 0.0007	397 12.02 $\pm$ 0.78	397 0.7484 $\pm$ 0.0076	397 0.8568 $\pm$ 0.0051

398 QT (QTcF), ST-segment deviation at J+60 ms (ST@J+60), and P- and T-wave durations (P dur, T  
399 dur), and compute record-level median absolute errors between generated and real ECGs.

400 Tables 1 report MAE, NRMSE, MAE<sub>HR</sub>, rCLIP, rFID on MIMIC-IV-ECG and PTB-XL dataset.  
401 Across both datasets, SE-Diff consistently outperforms SSDM, WGAN, BeatDiff, and DiffuSETS  
402 on all metrics, indicating improved waveform reconstruction, more accurate heart rate estimation,  
403 and stronger ECG–text alignment. The ablations (SE-Diff w/o Sim, w/o InterLead, w/o Exp) each  
404 degrade one or more metrics, highlighting the importance of simulator guidance, inter-lead con-  
405 straints, and experience-based conditioning.

407 Table 2: Beat-level morphology & interval fidelity (MAE) on MIMIC-IV-ECG and PTB-XL  
408 datasets. All values are record-level medians; lower is better ( $\downarrow$ ).

410 Model	410 PR $\downarrow$	410 QRSd $\downarrow$	410 QT $\downarrow$	410 QTcF $\downarrow$	410 ST@J+60 $\downarrow$	410 P dur $\downarrow$	410 T dur $\downarrow$
<b>MIMIC-IV-ECG (internal)</b>							
412 SSDM	413 7.70 $\pm$ 8	413 32.00 $\pm$ 1	413 25.11 $\pm$ 11	413 23.25 $\pm$ 11	413 0.16 $\pm$ 0	413 5.30 $\pm$ 3	413 23.71 $\pm$ 9
413 WGAN	414 18.01 $\pm$ 9	414 29.64 $\pm$ 9	414 21.02 $\pm$ 5	414 23.83 $\pm$ 2	414 0.13 $\pm$ 0	414 6.01 $\pm$ 0	414 24.01 $\pm$ 4
414 BeatDiff	415 10.34 $\pm$ 5	415 18.23 $\pm$ 6	415 14.23 $\pm$ 6	415 9.43 $\pm$ 5	415 0.17 $\pm$ 0	415 7.00 $\pm$ 1	415 30.02 $\pm$ 3
415 DiffuSETS	416 14.81 $\pm$ 6	416 11.71 $\pm$ 7	416 8.20 $\pm$ 3	416 9.71 $\pm$ 4	416 0.04 $\pm$ 0	416 5.60 $\pm$ 1	416 8.71 $\pm$ 7
<b>SE-Diff</b>							
416 <b>SE-Diff</b>	417 <b>7.30 <math>\pm</math> 3</b>	417 <b>10.71 <math>\pm</math> 4</b>	417 <b>4.50 <math>\pm</math> 2</b>	417 <b>7.88 <math>\pm</math> 2</b>	417 <b>0.03 <math>\pm</math> 0</b>	417 <b>2.50 <math>\pm</math> 0</b>	417 <b>6.80 <math>\pm</math> 3</b>
417 w/o InterLead	418 11.11 $\pm$ 3	418 15.01 $\pm$ 5	418 11.91 $\pm$ 5	418 15.84 $\pm$ 7	418 0.03 $\pm$ 0	418 4.00 $\pm$ 1	418 8.10 $\pm$ 4
418 w/o Exp	419 9.81 $\pm$ 9	419 14.11 $\pm$ 13	419 8.20 $\pm$ 3	419 12.61 $\pm$ 3	419 0.04 $\pm$ 0	419 4.00 $\pm$ 1	419 6.90 $\pm$ 3
419 w/o Sim	420 12.11 $\pm$ 7	420 13.21 $\pm$ 6	420 8.40 $\pm$ 4	420 8.89 $\pm$ 7	420 0.04 $\pm$ 0	420 5.10 $\pm$ 1	420 13.11 $\pm$ 1
<b>PTB-XL (external)</b>							
420 SSDM	421 13.71 $\pm$ 8	421 13.61 $\pm$ 8	421 30.22 $\pm$ 9	421 27.69 $\pm$ 14	421 0.33 $\pm$ 0	421 12.71 $\pm$ 5	421 19.11 $\pm$ 10
421 WGAN	422 14.01 $\pm$ 7	422 17.45 $\pm$ 13	422 12.01 $\pm$ 10	422 12.97 $\pm$ 10	422 0.18 $\pm$ 0	422 7.33 $\pm$ 3	422 25.68 $\pm$ 13
422 BeatDiff	423 13.41 $\pm$ 4	423 17.40 $\pm$ 6	423 9.81 $\pm$ 3	423 8.94 $\pm$ 2	423 0.73 $\pm$ 1	423 4.10 $\pm$ 1	423 18.81 $\pm$ 6
423 DiffuSETS	424 7.70 $\pm$ 2	424 12.61 $\pm$ 7	424 11.51 $\pm$ 4	424 13.75 $\pm$ 5	424 0.10 $\pm$ 0	424 9.11 $\pm$ 6	424 13.20 $\pm$ 5
<b>SE-Diff</b>							
424 <b>SE-Diff</b>	425 <b>3.90 <math>\pm</math> 1</b>	425 <b>10.01 <math>\pm</math> 2</b>	425 <b>5.20 <math>\pm</math> 1</b>	425 <b>8.68 <math>\pm</math> 2</b>	425 <b>0.07 <math>\pm</math> 0</b>	425 <b>3.00 <math>\pm</math> 1</b>	425 <b>9.31 <math>\pm</math> 8</b>
425 w/o InterLead	426 7.00 $\pm$ 1	426 10.31 $\pm$ 5	426 8.40 $\pm$ 4	426 10.91 $\pm$ 3	426 0.07 $\pm$ 0	426 4.40 $\pm$ 1	426 13.61 $\pm$ 8
426 w/o Sim	427 9.21 $\pm$ 4	427 12.91 $\pm$ 5	427 12.51 $\pm$ 4	427 14.20 $\pm$ 2	427 0.07 $\pm$ 0	427 4.50 $\pm$ 1	427 12.31 $\pm$ 3

427 Tables 2 summarize the beat-level morphology and interval metrics on MIMIC-IV-ECG and PTB-  
428 XL. SE-Diff achieves the lowest median errors across all intervals (PR, QRSd, QT/QTcF, ST@J+60,  
429 P dur, T dur), demonstrating that it not only matches global signal statistics but also preserves  
430 beat-level timing and morphology more faithfully than competing models. Consistent performance  
431 gaps between SE-Diff and its ablations further support that the simulator- and experience-enhanced  
432 design improves clinically relevant morphology and interval fidelity.

432 4.2 DOWNSTREAM ECG CLASSIFICATION  
433

434 We evaluate whether SE-Diff mitigates severe class imbalance in downstream ECG classification by  
435 augmenting minority classes with model-generated ECGs. Training distributions are intentionally  
436 skewed; evaluation uses a fixed, class-balanced test set. We compare four regimes: (i) *Unbalanced*  
437 (real-only, skewed), (ii) *Balanced* (real-only, fully balanced; reference upper bound), and (iii) *sev-*  
438 *eral synthetic augmentation strategies that oversample the minority class using SSDM, WGAN,*  
439 *BeatDiff, DiffuSETS, or SE-Diff. We report F1, accuracy (Acc.), and AUROC (AUC).*

440 **Imbalanced Gender Classification.** We train a binary classifier to predict sex (Female vs. Male).  
441 As shown in Table 3, all generative augmentation methods improve over the *Unbalanced* baseline,  
442 but SE-Diff yields the largest gains, raising F1 from 42% to 58% and AUC from 46% to 58%. This  
443 substantially narrows the gap to the *Balanced* upper bound (F1/AUC 62%), and SE-Diff consistently  
444 outperforms SSDM, WGAN, BeatDiff, and DiffuSETS under the same skew.

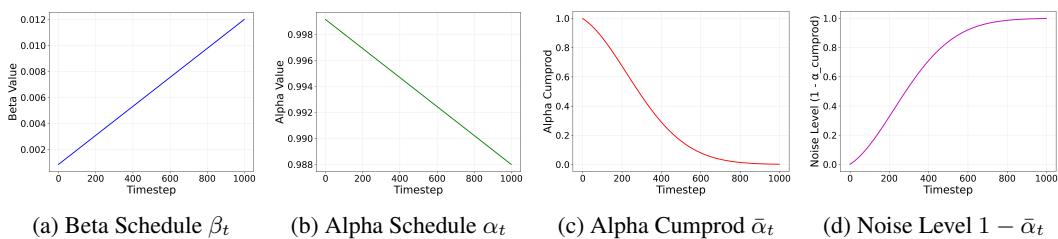
445 **Rare-disease Classification.** We train a classifier to distinguish Sinus rhythm from supraventricular  
446 tachycardia (SVT), treating SVT as the minority class. Table 3 shows larger relative gains on SVT,  
447 indicating that synthetic augmentation is particularly effective when physiological heterogeneity is  
448 high and labeled minority examples are scarce. SE-Diff recovers a substantial fraction of this gap,  
449 especially on the minority SVT class.

450 Table 3: Downstream ECG classification under severe class imbalance.  
451

452 <b>Model</b>	453 <b>Male = 10% Female</b>			454 <b>SVT = 10% Sinus</b>		
	455 F1 (%), ↑	456 Acc. (%), ↑	457 AUC (%), ↑	458 F1 (%), ↑	459 Acc. (%), ↑	460 AUC (%), ↑
454 SSDM	42 ± 0	54 ± 1	49 ± 2	56 ± 1	63 ± 2	81 ± 2
455 WGAN	42 ± 0	54 ± 1	49 ± 2	57 ± 2	63 ± 2	82 ± 2
456 BeatDiff	44 ± 2	55 ± 2	51 ± 2	60 ± 2	67 ± 2	84 ± 3
457 DiffuSETS	44 ± 3	54 ± 2	54 ± 1	70 ± 1	68 ± 1	84 ± 2
458 SE-Diff (ours)	<b>58 ± 1</b>	<b>58 ± 1</b>	<b>58 ± 2</b>	<b>72 ± 2</b>	<b>71 ± 0</b>	<b>85 ± 2</b>
459 Unbalanced	42 ± 0	54 ± 1	46 ± 2	56 ± 1	62 ± 0	80 ± 1
460 Balanced	62 ± 0	62 ± 0	62 ± 1	79 ± 1	80 ± 2	93 ± 1

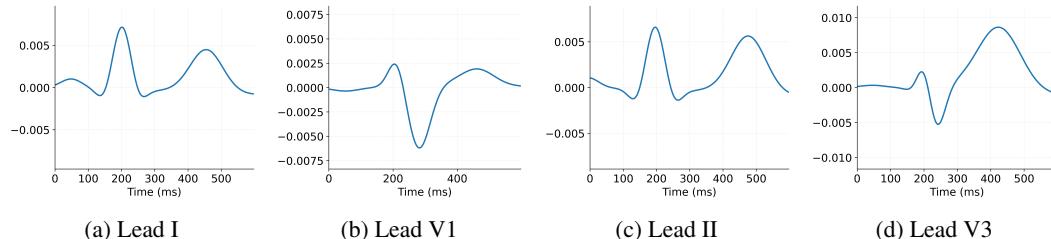
461 4.3 MECHANISTIC ANALYSIS OF SE-DIFF  
462

463 **Noise Scheduling Analysis.** Figure 2 summarizes the forward process under our linear noise sched-  
464 ule. Panel (a) shows the per-step noise increment increasing steadily, while (b) displays the corre-  
465 sponding signal retention factor decreasing slightly each step. The cumulative signal fraction in (c)  
466 drops smoothly from near one to near zero, and the accumulated noise in (d) rises monotonically  
467 and saturates in late timesteps. This profile yields a gradual, well-conditioned reverse trajectory:  
468 early steps recover global rhythm and cross-lead coherence, and later steps refine P/QRS/T mor-  
469 phology and suppress residual artifacts. We therefore adopt this schedule for SE-Diff as it offers an  
470 interpretable progression and stable behavior across sampling budgets.  
471

472 Figure 2: Noise scheduling analysis showing the progression of noise and signal factors throughout  
473 the diffusion process.  
474

475 **Case Study for ECG Simulator.** To visualize the simulator’s morphology, Fig. 3 presents four  
476 random single-cycle templates (one lead per label): (A) *Sinus rhythm, Lead I*. Upright P wave, nar-  
477 row QRS complex, and concordant T wave provide a clean normal reference for comparison. (B)  
478 *Ventricular pacing, Lead VI*. A wide, predominantly negative QRS complex (QS/deep S), reflect-  
479 ing pacing/LBBB-like activation, clearly departs from normal conduction. (C) *Sinus rhythm with*

486 *first-degree AV block, Lead II.* A P wave followed by an elongated isoelectric segment before the  
 487 QRS complex qualitatively indicates PR-interval prolongation. *(D) Consider acute ST-elevation MI,*  
 488 *Lead V3.* Convex ST-segment elevation after the J point is characteristic of anteroseptal involve-  
 489 ment. The simulator serves as a morphology prior and qualitative oracle within SE-Diff, enabling  
 490 clear visual audits and morphology-aware ablations without confounding rhythm variability, and  
 491 providing guidance to the diffusion model.



500 Figure 3: Representative single-cycle ECG waveforms generated from our simulator. Panel A: sinus  
 501 rhythm (Lead I). Panel B: ventricular pacing (Lead V1). Panel C: sinus rhythm with first-degree AV  
 502 block (Lead II). Panel D: consider acute ST-elevation MI (Lead V3).

503 **Case Study for ECG Generation.** Figure 4 compares a 10 s, 12-lead ECG generated by SE-Diff  
 504 with its paired reference for a case conditioned on “*Sinus rhythm*” (male, 65 y, HR 94 bpm). The  
 505 generation preserves canonical sinus morphology—each *P* wave preceding a narrow QRS com-  
 506 plex with an appropriate PR interval—and shows coherent R-wave progression across the precordial  
 507 leads, with R-R intervals consistent with the target rate. Clinically, the SE-Diff tracing appears  
 508 cleaner than the reference: baseline wander and high-frequency artifacts are attenuated, yielding  
 509 crisper ST segments and T-wave contours without distorting morphology. This qualitative finding  
 510 aligns with the model’s design: simulator-informed constraints and experience-augmented  
 511 conditioning steer the diffusion process toward physiologically plausible, low-noise signals.



525 (a) SE-Diff generated ECG. (b) Real ECG.

526 Figure 4: Case Study for ECG Generation.

## 528 5 CONCLUSION

530 We introduced SE-Diff, a conditional latent-diffusion framework for 10s, 12-lead ECG syn-  
 531 thesis that couples a VAE latent space with a Beat Decoder and simulator-informed regularizers, and  
 532 strengthens conditioning via experience retrieval from EHRs. Across benchmarks, SE-Diff improves  
 533 signal fidelity, preserves inter-lead physiology, and achieves tighter diagnostic/semantic alignment,  
 534 while also enhancing downstream classification when used for data augmentation. Ablations  
 535 confirm that both the ODE-based guidance (Euler and inter-lead constraints) and retrieval-augmented  
 536 conditioning contribute materially to performance. Future work will extend SE-Diff to more clin-  
 537 ically meaningful applications (e.g., arrhythmia risk stratification, therapy response modeling, and  
 538 long-term ambulatory ECG) and evaluate robustness across institutions and rare presentations. We  
 539 believe SE-Diff offers a principled step toward physiologically grounded, clinically aligned genera-  
 tive modeling of ECGs.

540 6 ETHICS STATEMENT  
541

542 We adhere to the ICLR Code of Ethics. This study uses only de-identified data from MIMIC-IV-  
543 ECG (Gow et al., 2023; Johnson et al., 2023) and associated de-identified records from MIMIC-  
544 IV-CLINICAL (Johnson et al., 2023) under the applicable Data Use Agreements and credentialed-  
545 access requirements. No direct interaction with human subjects occurred; no personally identifiable  
546 information (PII/PHI) was accessed or released, and we made no attempt at re-identification. Heart-  
547 rate estimation (WFDB XQRS), resampling, and VAE-based encoding were performed on the de-  
548 identified waveforms; EHR linkage relied only on the dataset’s de-identified subject and admission  
549 keys within documented admission windows. This work is for research purposes only and is not a  
550 medical device. No conflicts of interest or sensitive sponsorships are present.

551  
552 7 REPRODUCIBILITY STATEMENT  
553

554 All information necessary to reproduce our results is documented in Appendix D, including dataset  
555 curation and preprocessing, model architectures, training and inference hyperparameters and sched-  
556 ules, and so on.

557  
558 REFERENCES  
559

560 Elmira Adib, Fatemeh Afghah, and John J. Prevost. Arrhythmia classification using cgan-augmented  
561 ecg signals. In *2022 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*,  
562 pp. 1865–1872. IEEE, 2022.

563 Valtino X Afonso, Willis J Tompkins, Truong Q Nguyen, and Shen Luo. Ecg beat detection using  
564 filter banks. *IEEE transactions on biomedical engineering*, 46(2):192–202, 1999.

565 Lisa Bedin and Coauthors. Leveraging an ecg beat diffusion model for morphological reconstruction  
566 from indirect signals. In *Advances in Neural Information Processing Systems*, volume 37, pp.  
567 84409–84446, 2024.

568 Vitalii Bondar, Serhii Semenov, Vira Babenko, and Dmytro Holovniak. Flowecg: Using flow match-  
569 ing to create a more efficient ecg signal generator. *arXiv preprint arXiv:2509.10491*, 2025. URL  
570 <https://arxiv.org/abs/2509.10491>. To appear in *Sensors, Devices and Systems –*  
571 *2025 Proceedings*.

572 John Charles Butcher and JC Butcher. *The numerical analysis of ordinary differential equations:*  
573 *Runge-Kutta and general linear methods*, volume 512. Wiley New York, 1987.

574 Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David Duvenaud. Neural ordinary dif-  
575 ferential equations. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

576 Philip De Chazal, Maria O’Dwyer, and Richard B Reilly. Automatic classification of heartbeats  
577 using ecg morphology and heartbeat interval features. *IEEE transactions on biomedical engi-*  
578 *neering*, 51(7):1196–1206, 2004.

579 Tomer Golany, Gal Lavee, Shai Tejman Yarden, and Kira Radinsky. Improving ecg classification  
580 using generative adversarial networks. In *Proceedings of the AAAI Conference on Artificial Intel-*  
581 *ligence*, 2020a.

582 Tomer Golany, Kira Radinsky, and Daniel Freedman. Simgans: Simulator-based generative adver-  
583 sarial networks for ecg synthesis to improve deep ecg classification. In *International Conference*  
584 *on Machine Learning*, pp. 3597–3606. PMLR, 2020b.

585 Tomer Golany, Daniel Freedman, and Kira Radinsky. ECG ODE-GAN: Learning ordinary differen-  
586 tial equations of ECG dynamics via generative adversarial learning. In *Proceedings of the AAAI*  
587 *Conference on Artificial Intelligence*, volume 35, pp. 134–141, 2021. doi: 10.1609/aaai.v35i1.  
588 16086.

594 Ary L. Goldberger, Luis A. N. Amaral, Leon Glass, Jeffrey M. Hausdorff, Plamen Ch. Ivanov,  
 595 Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and H. Eugene Stanley.  
 596 Physiobank, physiotookit, and physionet: Components of a new research resource for complex  
 597 physiologic signals. *Circulation*, 101(23):e215–e220, 2000. doi: 10.1161/01.CIR.101.23.e215.  
 598

599 Brian Gow, Benjamin Moody, Lucas Bulgarelli, Alistair E. W. Johnson, Li-wei H. Lehman, and  
 600 Roger G. Mark. Mimic-iv-ecg: A publicly available database of real-world 12-lead ecgs. Phys-  
 601 ioNet, 2023.

602 Tae-Seong Han, Jae-Wook Heo, Hakseung Kim, Cheol-Hui Lee, Hyub Huh, Eue-Keun Choi,  
 603 Hye Jin Kim, and Dong-Joo Kim. Diffusion-based electrocardiography noise quantification via  
 604 anomaly detection. *arXiv preprint arXiv:2506.11815*, 2025.

605

606 Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Baining  
 607 Guo. Efficient diffusion training via min-snr weighting strategy. In *Proceedings of the IEEE/CVF*  
 608 *international conference on computer vision*, pp. 7441–7451, 2023.

609

610 Awni Y Hannun, Pranav Rajpurkar, Masoumeh Haghpanahi, Geoffrey H Tison, Cody Bourn,  
 611 Mintu P Turakhia, and Andrew Y Ng. Cardiologist-level arrhythmia detection and classifica-  
 612 tion in ambulatory electrocardiograms using a deep neural network. *Nature Medicine*, 25(1):  
 613 65–69, 2019. doi: 10.1038/s41591-018-0268-3.

614 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance, 2022.

615

616 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances*  
 617 *in Neural Information Processing Systems (NeurIPS)*, 2020.

618

619 Alistair E. W. Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng,  
 620 Tom J. Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, Li-wei H. Lehman, Leo A. Celi, and  
 621 Roger G. Mark. Mimic-iv, a freely accessible electronic health record dataset. *Scientific Data*, 10  
 622 (1):1, 2023. doi: 10.1038/s41597-022-01899-x.

623

624 Mohammad Kachuee, Shayan Fazeli, and Majid Sarrafzadeh. Ecg heartbeat classification: A deep  
 625 transferable representation. In *2018 IEEE International Conference on Healthcare Informatics*  
 (ICHI), pp. 443–444. IEEE, 2018.

626

627 Paul Kligfield, Leonard S Gettes, James J Bailey, Rory Childers, Barbara J Deal, E. William  
 628 Hancock, Gerard van Herpen, Jan A Kors, Peter Macfarlane, David M Mirvis, and et al.  
 629 Recommendations for the standardization and interpretation of the electrocardiogram: Part  
 630 i: The electrocardiogram and its technology. *Circulation*, 115(10):1306–1324, 2007. doi:  
 631 10.1161/CIRCULATIONAHA.106.180200.

632

633 Yongfan Lai, Jiabo Chen, Qinghao Zhao, Deyun Zhang, Yue Wang, Shijia Geng, Hongyan Li, and  
 634 Shenda Hong. Diffusets: 12-lead ecg generation conditioned on clinical text reports and patient-  
 635 specific information. *Patterns*, 2025a.

636

637 Yongfan Lai, Bo Liu, Xinyan Guan, Qinghao Zhao, Hongyan Li, and Shenda Hong. Ecgtwin: Per-  
 638 sonalized ecg generation using controllable diffusion model. *arXiv preprint arXiv:2508.02720*,  
 639 2025b.

640

641 Byeong Tak Lee, Yong-Yeon Jo, and Joon-Myoung Kwon. Optimizing neural network scale for ecg  
 642 classification. *arXiv preprint arXiv:2308.12492*, 2023.

643

644 Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal,  
 645 Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe  
 646 Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances in Neural*  
 647 *Information Processing Systems (NeurIPS)*, volume 33, pp. 9459–9474, 2020.

648

649 Dawei Li, Tienan Chen, Cong Liu, Liqi Liao, Shaoping Chen, Yongqiang Cui, Li Zhang, Yiyu Shi,  
 650 and Xiaowei Xu. Semi-supervised gan for enhancing electrocardiogram time series diagnostics.  
 651 *Biomedical Signal Processing and Control*, 110:108058, 2025.

648 Juan Miguel Lopez Alcaraz and Nils Strodthoff. Diffusion-based conditional ECG generation with  
 649 structured state space models. *Computers in Biology and Medicine*, 163:107115, 2023. doi:  
 650 10.1016/j.combiomed.2023.107115.

651 Juan Miguel López Alcaraz and Nils Strodthoff. Diffusion-based conditional ecg generation with  
 652 structured state space models. *Computers in Biology and Medicine*, 163:107115, 2023.

653

654 Marek Malik. Heart rate variability: Standards of measurement, physiological interpretation, and  
 655 clinical use: Task force of the european society of cardiology and the north american society for  
 656 pacing and electrophysiology. *Annals of Noninvasive Electrocardiology*, 1(2):151–181, 1996.

657 Marek Malik and A John Camm. Heart rate variability. *Clinical cardiology*, 13(8):570–576, 1990.

658

659 Samira Mavaddati. Ecg arrhythmias classification based on deep learning methods and transfer  
 660 learning technique. *Biomedical Signal Processing and Control*, 101:107236, 2025.

661

662 Patrick E McSharry, Gari D Clifford, Lionel Tarassenko, and Leonard A Smith. A dynamical model  
 663 for generating synthetic electrocardiogram signals. *IEEE transactions on biomedical engineering*,  
 50(3):289–294, 2003.

664

665 Louis Melville Milne-Thomson. *The calculus of finite differences*. American Mathematical Soc.,  
 666 2000.

667 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wier-  
 668 stra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint*  
 669 *arXiv:1312.5602*, 2013.

670

671 Nasser M Nasrabadi. Pattern recognition and machine learning. *Journal of electronic imaging*, 16  
 672 (4):049901, 2007.

673

674 Nour Neifar, Achraf Ben-Hamadou, Afef Mdhaffar, and Mohamed Jmaiel. DiffECG: A versatile  
 probabilistic diffusion model for ECG signals synthesis, 2023. Accepted at IEEE SERA 2024.

675

676 Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models.  
 677 In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, volume 139  
 of *Proceedings of Machine Learning Research*, pp. 8162–8171. PMLR, 2021.

678

679 Justice Ou, Tinglin Huang, Yilun Zhao, Ziyang Yu, Peiqing Lu, and Rex Ying. Experience retrieval-  
 680 augmentation with electronic health records enables accurate discharge qa. *arXiv preprint*  
 681 *arXiv:2503.17933*, 2025.

682

683 Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual  
 reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial  
 684 Intelligence*, volume 32, 2018.

685

686 Christopher Rackauckas, Yingbo Ma, Julius Martensen, Collin Warner, Kirill Zubov, Rohit Supekar,  
 687 Dominic Skinner, and Ali Ramadhan. Universal differential equations for scientific machine  
 learning. *Proceedings of the National Academy of Sciences*, 118(46):e210702118, 2021. doi:  
 688 10.1073/pnas.210702118.

689

690 Pranav Rajpurkar, Awni Y. Hannun, Masoumeh Haghpanahi, Codie Bourn, and Andrew Y. Ng.  
 691 Cardiologist-level arrhythmia detection with convolutional neural networks. *CoRR*, 2017.

692

693 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Con-  
 694 ference on Computer Vision and Pattern Recognition (CVPR)*, pp. 10684–10695, 2022.

695

696 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomed-  
 697 ical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention –  
 698 MICCAI 2015*, volume 9351 of *Lecture Notes in Computer Science*, pp. 234–241. Springer, 2015.  
 699 doi: 10.1007/978-3-319-24574-4\_28.

700

701 Gregory A Roth, George A Mensah, Catherine O Johnson, and et al. Global burden of cardiovascular  
 diseases and risk factors, 1990–2019. *Journal of the American College of Cardiology*, 76(25):  
 2982–3021, 2020. doi: 10.1016/j.jacc.2020.11.010.

702 Yuling Sang, Abhirup Banerjee, Marcel Beetz, and Vicente Grau. Deep conditional generative model  
 703 for personalization of 12-lead electrocardiograms and cardiovascular risk prediction. *Frontiers in*  
 704 *Digital Health*, 7:1558589, 2025.

705 Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Joshua Susskind, Wenda Wang, and Russell Webb.  
 706 Learning from simulated and unsupervised images through adversarial training. In *Proceedings*  
 707 *of the IEEE conference on computer vision and pattern recognition*, pp. 2107–2116, 2017.

708 Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsuper-  
 709 vised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International*  
 710 *Conference on Machine Learning (ICML)*, volume 37 of *Proceedings of Machine Learning Re-*  
 711 *search*, pp. 2256–2265. PMLR, 2015.

712 Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben  
 713 Poole. Score-based generative modeling through stochastic differential equations. In *Interna-*  
 714 *tional Conference on Learning Representations (ICLR)*, 2021.

715 Leif Sörnmo and Pablo Laguna. *Bioelectrical signal processing in cardiac and neurological appli-*  
 716 *cations*, volume 8. Academic Press, 2005.

717 Endre Süli and David F Mayers. *An introduction to numerical analysis*. Cambridge university press,  
 718 2003.

719 Task Force of the European Society of Cardiology and the North American Society of Pacing and  
 720 Electrophysiology. Heart rate variability: Standards of measurement, physiological interpretation,  
 721 and clinical use. *Circulation*, 93(5):1043–1065, 1996. doi: 10.1161/01.CIR.93.5.1043.

722 Vineeth Thambawita, Jonas L Isaksen, Steven Hicks, Jonas Ghouse, Gustav Ahlberg, Allan Lin-  
 723 neberg, Niels Grarup, Christina Ellervik, Morten S Olesen, Tobias Hansen, Claus Graff, et al.  
 724 Deepfake electrocardiograms using generative adversarial networks are the beginning of the  
 725 end for privacy issues in medicine. *Scientific Reports*, 11(1):21896, 2021. doi: 10.1038/s41598-021-01295-2.

726 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,  
 727 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Infor-*  
 728 *mation Processing Systems (NeurIPS)*, volume 30, 2017.

729 Patrick Wagner, Nils Strodtthoff, Ralf-Dieter Bousseljot, David Kreiseler, Felix I Lunze, Wojciech  
 730 Samek, and Tobias Schaeffter. PTB-XL, a large publicly available electrocardiography dataset.  
 731 *Scientific Data*, 7(1):154, 2020. doi: 10.1038/s41597-020-0495-6.

732 Chuanjiang Wang, Junhao Ma, Guohui Wei, and Xiujuan Sun. Analysis of cardiac arrhythmias  
 733 based on resnet-icbam-2dcnn dual-channel feature fusion. *Sensors*, 25(3):661, 2025.

734 Dean Zadok, Tom Hirshberg, Amir Biran, Kira Radinsky, and Ashish Kapoor. Explorations and  
 735 lessons learned in building an autonomous formula sae car from simulations. *SIMULTECH*,  
 736 2019.

737 Md Haider Zama and Friedhelm Schwenker. Ecg synthesis via diffusion-based state space aug-  
 738 mented transformer. *Sensors*, 23(19):8328, 2023. doi: 10.3390/s23198328.

739 Beatrice Zanchi, Giuliana Monachino, Luigi Fiorillo, Giulio Conte, Angelo Auricchio, Athina Tzo-  
 740 varia, and Francesca D Faraci. Synthetic ecg signals generation: A scoping review. *Computers in*  
 741 *Biology and Medicine*, 184:109453, 2025.

742

743

744

745

746

747

748

749

750

751

752

753

754

755

## A USE OF LARGE LANGUAGE MODELS

We used a large language model (LLM) solely for writing assistance—specifically, to correct grammar, improve wording, and enhance clarity. The LLM did not contribute to research design, data analysis, modeling, experiments, or interpretation of results. All technical content and conclusions were authored and verified by the authors, who take full responsibility for the manuscript.

756 **B RELATED WORK**  
757758 **B.1 GENERATIVE MODELS FOR ECG**  
759

760 Simulation has repeatedly improved data efficiency in sequential decision making—both in imitation  
761 learning and in reinforcement learning—by narrowing the gap between training and deployment  
762 (Zadok et al., 2019; Mnih et al., 2013). In parallel, generative models have been used to expand training  
763 corpora: in vision, SimGAN refines synthetic images with unlabeled real data (Shrivastava et al.,  
764 2017); in cardiology, augmenting classifiers with GAN-generated heartbeats improves performance  
765 (Golany et al., 2020a). Beyond beat-level augmentation, adversarial models have produced realistic  
766 multi-lead “DeepFake” ECGs for privacy and data scarcity mitigation (Thambawita et al., 2021),  
767 and mechanism-aware variants embed ordinary differential equations to better capture depolariza-  
768 tion–repolarization dynamics (Golany et al., 2021). However, GANs can be unstable and prone  
769 to mode collapse in multi-lead, multi-label regimes. Denoising diffusion and score-based models  
770 offer a likelihood-grounded alternative with strong mode coverage and stable training (Ho et al.,  
771 2020; Song et al., 2021). Recent ECG adaptations include conditional diffusion with structured  
772 state-space backbones (SSSD-ECG) (Lopez Alcaraz & Strodthoff, 2023), generalized diffusion for  
773 generation/imputation/forecasting (Neifar et al., 2023), state-space/transformer hybrids (Zama &  
774 Schwenker, 2023), and text/metadata-conditioned synthesis (DiffuSETS) (Lai et al., 2025a). The  
775 field is also trending toward personalization and physiological consistency: conditional models in-  
776 incorporate patient metadata or anatomy to produce more plausible 12-lead signals (Sang et al., 2025),  
777 and diffusion frameworks create patient “digital twins” (Lai et al., 2025b). Hybrid uses couple gener-  
778 ative modeling with signal-quality assessment and anomaly detection (Han et al., 2025), while  
779 semi-supervised GANs aim to better capture temporal dynamics (Li et al., 2025). Very recent work  
780 explores flow-matching as a faster alternative to iterative diffusion for ECG synthesis, reducing  
sampling cost while targeting comparable fidelity (Bondar et al., 2025).

781 **B.2 PHYSIOLOGICAL ECG SIMULATORS**  
782

783 Compact physiological simulators capture stereotyped P–QRS–T morphology with low-dimensional  
784 differential equations. The canonical ECGSYN model uses a three-dimensional limit-cycle oscil-  
785 lator whose phase-locked Gaussian components generate P, QRS, and T deflections, while stochas-  
786 tic control of instantaneous heart rate reproduces realistic RR patterns and HRV statistics (e.g.,  
787 mean/SD of RR, low- and high-frequency spectral peaks) (McSharry et al., 2003; Task Force of the  
788 European Society of Cardiology and the North American Society of Pacing and Electrophysiology,  
789 1996). Open implementations (e.g., PhysioNet ECGSYN) enable reproducible waveform synthe-  
790 sis and stress-testing (Goldberger et al., 2000). However, globally fixed morphology templates and  
791 linear lead projections limit expressivity under rhythm changes, conduction abnormalities, and non-  
792 stationary repolarization. Hybrid approaches mitigate these issues by coupling mechanistic priors  
793 with learnable components—via neural ODEs or universal differential equations—to preserve phys-  
794 ical structure while fitting data; conditioning on anatomy further improves inter-lead realism (Chen  
795 et al., 2018; Rackauckas et al., 2021; Sang et al., 2025).

796 **B.3 ECG CLASSIFICATION**  
797

798 Classical ECG pipelines segment signals into beats with robust QRS detectors (e.g., Pan–Tompkins;  
799 Afonso et al.) and derive interval/morphology descriptors before applying shallow classifiers such  
800 as linear discriminants or SVMs (Afonso et al., 1999; De Chazal et al., 2004; Nasrabadi, 2007).  
801 With deep learning, end-to-end models on raw waveforms supplanted hand-crafted features and  
reached cardiologist-level performance in single- and ambulatory-lead arrhythmia detection (Ra-  
802 jpurkar et al., 2017; Hannun et al., 2019); at the beat level, residual CNNs are particularly effective,  
803 and large multi-lead corpora such as PTB-XL have enabled high-capacity models and rigorous multi-  
804 label benchmarking (Kachuee et al., 2018; Wagner et al., 2020). Recent work refines architectures  
805 and training regimes—dual-channel networks that fuse ResNet-ICBAM with 2D-CNN features em-  
806 phasize region-of-interest cues (Wang et al., 2025), ECG-specific scaling laws suggest shallower but  
807 wider networks outperform vision-oriented designs (Lee et al., 2023), and transfer learning on trans-  
808 formed signals improves performance under class imbalance (Mavaddati, 2025). Synthetic data from  
809 GANs and diffusion models is now routinely used for augmentation, with semi-supervised variants  
further boosting diagnostic accuracy (Li et al., 2025). In our experiments, we adopt a strong ResNet

heartbeat classifier and evaluate whether simulator-enhanced diffusion synthesis (SE-Diff) improves generalization under class imbalance and limited labels by augmenting training with physiologically plausible, label-consistent synthetic beats.

## C ECG SIMULATOR CALIBRATION WITH STABILIZERS AND MORPHOLOGY PRIORS

**Motivation.** Naive least-squares calibration of the ODE-based simulator (Sec. 2.2) often fits the sharp QRS complex yet drifts later in the window and may flip polarity. The main causes are small baseline trends and scale mismatches between simulated and observed signals, and an under-constrained morphology (especially T-wave width). We introduce lightweight, differentiable stabilizers that improve convergence and yield physiologically plausible parameters without altering simulator dynamics.

**Trend-aware alignment and fidelity.** Let  $z_\eta(t)$  denote the simulated voltage with parameters  $\eta$ . Rather than compare  $z_\eta$  directly to the observation  $y(t)$ , align via a three-parameter affine–trend model with scale  $s$ , offset  $c$ , and linear trend  $b$ :

$$\hat{y}(t) = c + s z_\eta(t) + b(t - \bar{t}), \quad \bar{t} = \frac{1}{T} \sum_{t=1}^T t. \quad (27)$$

At each iteration  $(c, s, b)$  are obtained by least squares and are differentiable in  $z_\eta$ . The fidelity term is

$$\mathcal{L}_{\text{mse}}(\eta) = \frac{1}{T} \sum_{t=1}^T (\hat{y}(t) - y(t))^2, \quad (28)$$

augmented by a small scale regularizer to prevent rare runaway gains:

$$\mathcal{L}_s(\eta) = \lambda_s s^2, \quad \lambda_s \in [10^{-6}, 10^{-5}]. \quad (29)$$

**Morphology priors (widths and amplitudes).** Using the McSharry parameterization, each deflection  $\beta \in \{P, Q, R, S, T\}$  has amplitude  $a_\beta$ , width  $b_\beta > 0$ , and phase  $\theta_\beta$ . Enforce positivity with  $b_\beta = \text{softplus}(\tilde{b}_\beta) + \varepsilon$  ( $\varepsilon = 10^{-3}$ ) and shrink widths toward physiological targets  $b_\beta^*$ :

$$\mathcal{L}_{\text{width}}(\eta) = \lambda_b \sum_{\beta} w_\beta (b_\beta - b_\beta^*)^2, \quad (b_P^*, b_Q^*, b_R^*, b_S^*, b_T^*) = (0.20, 0.08, 0.10, 0.08, 0.32), \quad (30)$$

with  $w_T = 2$  and  $w_\beta = 1$  otherwise to prevent absorbing baseline drift via an overly broad T wave. A mild amplitude penalty

$$\mathcal{L}_{\text{amp}}(\eta) = \lambda_a \sum_{\beta} a_\beta^2 \quad (31)$$

discourages attributing variability solely to the global scale  $s$  in Eq. 27.

**Phase ordering.** To preserve the physiological ordering of  $\{P, Q, R, S, T\}$  on the unit circle, introduce a global phase shift  $\Delta\theta$  and wrap phases as  $\theta_\beta \leftarrow \text{wrap}(\tilde{\theta}_\beta + \Delta\theta)$ . A hinge penalty with margin  $m$  enforces monotonicity:

$$\mathcal{L}_{\text{ord}}(\eta) = \lambda_{\text{ord}} \sum_{i=1}^4 \max\{0, \theta_i - \theta_{i+1} + m\}, \quad m \approx 0.05 \text{ rad.} \quad (32)$$

This term typically decays after a few epochs and can be disabled once ordering stabilizes.

**Objective and optimization.** The calibration loss is

$$\mathcal{L}(\eta) = \mathcal{L}_{\text{mse}} + \mathcal{L}_s + \mathcal{L}_{\text{width}} + \mathcal{L}_{\text{amp}} + \mathcal{L}_{\text{ord}}. \quad (33)$$

Optimize with AdamW (cosine decay with warmup), followed by a brief L-BFGS refinement. The same Euler sub-stepping and burn-in used at inference are applied during training to maintain integrator consistency.

864 **Polarity canonicalization (post hoc).** Because lead inversions are common, canonicalize polarity  
 865 after fitting: re-simulate  $z_\eta$ , re-estimate  $\hat{y}(t) = c + s z_\eta(t)$  (no slope), and, if  $s < 0$ , flip all  
 866 amplitudes  $\{a_\beta\}$  once. This step is outside the loss and standardizes reported parameters.  
 867

868 **Default hyperparameters.** Unless specified otherwise, use  $\lambda_b = 5 \times 10^{-3}$  (with  $w_T=2$ ),  $\lambda_a =$   
 869  $4 \times 10^{-4}$ ,  $\lambda_{\text{ord}} = 10^{-4}$ ,  $\lambda_s = 10^{-6}$ ,  $m = 0.05$ , and  $\varepsilon = 10^{-3}$ . These values are deliberately  
 870 weak—sufficient to avoid the systematic drifts above without overriding the data.  
 871

872 Trend-aware alignment absorbs baseline wander, width priors (especially for  $T$ ) prevent compensatory  
 873 morphology stretching, the small scale penalty stabilizes gain, and the optional ordering term  
 874 removes rare phase crossings. Together, these stabilizers reduce late-window drift and polarity mis-  
 875 matches while preserving interpretable parameters  $\{a_\beta, b_\beta, \theta_\beta\}$ .  
 876

## 877 D IMPLEMENTATION DETAILS.

878 All models are trained in PyTorch with AMP on a single NVIDIA H200 using AdamW (lr  $1 \times 10^{-4}$   
 879 with cosine decay to  $1 \times 10^{-5}$ ), gradient clipping/accumulation (global batch 4096), for 200  
 880 epochs with early stopping. Diffusion uses  $T=1000$  steps, linear  $\beta_t \in [8.5 \times 10^{-4}, 1.2 \times 10^{-2}]$   
 881 (DDPMscheduler), and classifier-free guidance. The VAE has 4 latent channels; the en-  
 882 coder/decoder are multi-resolution with residual blocks, attention, and skip connections; training  
 883 uses  $\mathcal{L}_{\text{MSE}} + \mathcal{L}_{\text{KL}}$  ( $\lambda_{\text{KL}}=1$ ). A lightweight beat decoder predicts the first beat ( $L_c=300$  at 500Hz;  
 884 R-peaks via NeuroKit2). The denoiser is a 7-stage 1D U-Net (kernel 7) with self/cross-attention  
 885 (8 heads, width 16–64) consuming text embeddings (1536-d) plus metadata (age, sex, heart rate).  
 886 Physiology-aware training adds the Euler simulator loss ( $\lambda=3 \times 10^{-3}$ ) and the inter-lead constraint  
 887 ( $\gamma=5 \times 10^{-2}$ ); class-wise simulator parameters  $\eta$  are prefit from 200 beats/label.  
 888

889 We use MIMIC-IV-EKG with simplified rhythm labels. Free-text diagnostic reports are  
 890 cleaned and normalized, then mapped to a compact multi-label taxonomy (e.g., sinus  
 891 rhythm/brady/tachy, atrial fibrillation/flutter, PAC/PVC, bundle-branch block, LVH/RVH, prolonged  
 892 QT, ST/T abnormalities, ischemia/infarct). Texts are embedded with a pretrained text encoder  
 893 (text-embedding-ada-002). For simulator-informed diffusion, we pre-compute class-wise  
 894 simulator parameters for the top-20 EKG categories and use them during training. For EKG gen-  
 895 eration, we sample 100 waveforms per setting and compute metrics. For downstream EKG clas-  
 896 sification, we form balanced subsets with 200 samples per class and train a lightweight MLP that  
 897 flattens latents (512-d) and applies two fully connected layers (128→64) with BatchNorm, ReLU,  
 898 and Dropout (0.5), followed by a linear output; optimization uses cross-entropy with AdamW and  
 899 early stopping. In all experiments, we set  $\lambda_1 = \lambda_2 = \lambda_3 = 1$  and use  $k = 3$  retrieved admissions for  
 900 experience-based conditioning. And we employ the ‘gpt-4.1-mini’ model for generating the experi-  
 901 ence report  $r$ . The Beat Decoder is trained jointly with the VAE, and when R-peak detection fails or  
 902 no class-wise simulator parameters are available for a given sample, we simply skip the beat/spectral  
 903 or corresponding simulator losses for that sample to avoid instability while keeping the rest of the  
 904 training unchanged.  
 905

906 From an ethics and privacy perspective, all experiments are conducted on de-identified public  
 907 datasets (MIMIC-IV-EKG/EHR and PTB-XL), and we enforce strict *patient-level* splits: records  
 908 used at validation/test time are never included in the retrieval pool, so the EHR retrieval context  
 909 cannot contain labels or text from the same patient. Retrieval is restricted to the training cohort  
 910 only. Moreover, the simulator-informed and experience-augmented model maintains its advantage  
 911 on PTB-XL as an external cohort, suggesting that SE-Diff provides utility without simply memoriz-  
 912 ing individual training cases or overfitting to specific MIMIC records.  
 913

## 914 E PROMPT EXAMPLE

915 The prompt example in our SE-Diff can be shown in Figure 5.  
 916

## 917 F CONDITIONAL LATENT DIFFUSION—IMPLEMENTATION DETAILS

918 *Context construction.* The context  $c$  concatenates (i) token embeddings from the clinical report  
 919  $E_{\text{text}} \in \mathbb{R}^{m \times d_c}$  (from a frozen clinical text encoder) and (ii) a single metadata token  $e_{\text{meta}} \in \mathbb{R}^{1 \times d_c}$

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939

## == SYSTEM PROMPT ==

You are a cardiology-focused language model. Your job is to convert structured inputs into a concise, clinically accurate narrative suitable for a cardiology note. You must not invent findings and you must not provide treatment recommendations. You should prioritize the current patient's measurements and observations over any retrieved examples from similar patients.

The report you generate will serve as the conditioning text for a downstream text-to-ECG synthesis model. Therefore, the narrative must be internally consistent, and physiologically grounded so that it can guide waveform generation.

You will receive (1) current patient context, (2) the original machine generated diagnosis, and (3) retrieval diagnosis from clinically similar patients. Based on these inputs, you will produce exactly one paragraph of 60-120 words. Your language should be precise and use standard ECG terminology and units so the downstream generator can align text with waveform features.

## == DATA PROMPT ==

Original machine-generated diagnosis (most important first): Sinus tachycardia | Normal ECG except for rate  
Heart rate (bpm): 100  
Age (years): 52  
Gender: F  
Retrieved similar patients' diagnoses: Sinus rhythm | Normal ECG | Possible ectopic atrial tachycardia

Figure 5: Prompt examples.

formed by projecting age/sex/etc. to  $d_c$ . We set  $C = [E_{\text{text}}; e_{\text{meta}}] \in \mathbb{R}^{(m+1) \times d_c}$  and feed  $C$  to cross-attention at the bottleneck and (optionally) the two highest-resolution decoder blocks.

*U-Net blocks.* Each block consists of  $\text{Conv}(k=3) \rightarrow \text{GroupNorm} \rightarrow \text{SiLU} \rightarrow \text{Conv}(k=3)$  with a residual connection. Self-attention (multi-head) follows normalization at lower resolutions. Down-sampling uses stride-2 convolutions; up-sampling uses nearest-neighbor followed by  $\text{Conv}(k=3)$ . We use  $K$  resolution levels (e.g.,  $K=4$ ) and  $N$  blocks per level (e.g.,  $N=2$ ); channel width doubles on down-sampling and halves on up-sampling.

*Attention.* Cross-/self-attention are multi-head with head dimension  $d_a$  (e.g., 8 heads). Queries at the bottleneck attend to keys/values derived from  $C$ ; text tokens use sinusoidal positional encodings. Metadata is represented as a single learned token.

*Time embedding and FiLM.* The timestep embedding uses exponentially spaced sinusoids and a two-layer MLP (SiLU) to produce  $(\gamma_t, \beta_t)$  per block. A small MLP on  $\text{Pool}(C)$  (token average) yields  $(\gamma_c, \beta_c)$ . FiLM is applied after normalization to every residual block.

*Training setup.* We use a cosine noise schedule and optional learned variance (Nichol & Dhariwal, 2021), AdamW with EMA, and classifier-free guidance with unconditional dropout  $p_{\text{uncond}} \in [0.1, 0.2]$  and tuned guidance scale. Min-SNR- $\gamma$  weighting (Hang et al., 2023) is optional for stabilizing early and late timesteps. Simulator losses are computed on  $h = D_{\psi}^{\text{beat}}(z_0)$  and removed at inference.

## G ABLATION STUDY OF BEAT AND SPECTRAL LOSSES ON THE BEAT DECODER

To further clarify the relationship between the time-domain beat reconstruction loss  $L_{\text{beat}}$  in Eq. (13) and the spectral consistency loss  $L_{\text{spec}}$  in Eq. (15), we conduct an ablation study where we remove each term in turn while keeping the rest of the architecture and training protocol unchanged.  $L_{\text{beat}}$  encourages the Beat Decoder output  $h$  to match a single R-peak-centered cycle  $C(x)$  in the time domain, whereas  $L_{\text{spec}}$  matches the log-magnitude spectrum of  $h$  to the distribution of spectra computed over all detected beats  $\{C_j(x)\}_{j=1}^J$  in the same 10 s window. The two objectives therefore act on complementary aspects of the canonical beat:  $L_{\text{beat}}$  anchors  $h$  to an actual observed cycle, while  $L_{\text{spec}}$  regularizes its frequency content toward the ensemble of beats rather than the full 10 s signal.

Table 4 reports ECG generation performance on MIMIC-IV-ECG when using both losses (full SE-Diff), removing  $L_{\text{beat}}$  (“w/o Beat Loss”), and removing  $L_{\text{spec}}$  (“w/o Spec Loss”). Using both losses yields the best performance across all metrics. In particular, dropping  $L_{\text{beat}}$  leads to a drastic degradation in heart-rate accuracy ( $\text{MAE}_{\text{HR}}$  increases from 8.43 to 23.48) and noticeably worse rCLIP and rFID scores, indicating that the canonical beat is no longer well aligned with the temporal structure of the underlying rhythm. Removing  $L_{\text{spec}}$  instead mainly harms signal-level fidelity (higher NRMSE and weaker rCLIP/rFID), showing that spectral regularization is important for capturing realistic morphology and fine-scale waveform details.

These results empirically support that  $L_{\text{beat}}$  and  $L_{\text{spec}}$  are *complementary* rather than conflicting: the best-performing model is obtained when the canonical beat  $h$  is simultaneously constrained to resemble an actual beat in the time domain and to share spectral statistics with the set of beats present in the 10 s record.

Table 4: ECG generation performance on MIMIC-IV-ECG dataset.

Model	MAE $\downarrow$	NRMSE $\downarrow$	$\text{MAE}_{\text{HR}} \downarrow$	rCLIP Score $\uparrow$	rFID Score $\uparrow$
<b>SE-Diff (ours)</b>	<b><math>0.0923 \pm 0.0021</math></b>	<b><math>0.0714 \pm 0.0010</math></b>	<b><math>8.43 \pm 0.42</math></b>	<b><math>0.9470 \pm 0.0029</math></b>	<b><math>0.9509 \pm 0.0035</math></b>
w/o Beat Loss	$0.0951 \pm 0.0020$	$0.0743 \pm 0.0016$	$23.48 \pm 1.07$	$0.9393 \pm 0.0026$	$0.9140 \pm 0.0055$
w/o Spec Loss	$0.0942 \pm 0.0023$	$0.0797 \pm 0.0014$	$16.47 \pm 1.23$	$0.9257 \pm 0.0030$	$0.9073 \pm 0.0055$

## H INFERENCE EFFICIENCY OF SE-DIFF FOR 10 s, 12-LEAD ECG GENERATION

To show the inference efficiency, we report the wall-clock latency of SE-Diff for generating 10 s, 12-lead ECGs in Table 5. For each batch size, we measure the time required to run 1000 diffusion sampling steps and average the runtime over three batches.

The results show that the total time per batch remains nearly constant for batch sizes between 32 and 256, so the per-sample latency decreases from 0.425 s at batch size 32 to 0.050 s at batch size 256. Even for very large batches (up to 4096), the wall-clock time grows sublinearly with the batch size, and the per-sample latency can be reduced to 0.0205 s. This indicates that SE-Diff can efficiently exploit batching during inference, making large-scale ECG synthesis practical under our current implementation.

Table 5: Inference latency of the model for generating one 10 s, 12-lead ECG under different batch sizes. Times are averaged over 3 batches per configuration, with 1000 diffusion sampling steps.

Batch size	Time / batch (s)	Time / sample (s)
32	13.55	0.425
128	13.65	0.105
256	13.40	0.050
512	15.05	0.0295
1024	24.90	0.0245
2048	45.00	0.0220
4096	84.75	0.0205

## I EFFECTS OF NOISE SCHEDULE AND SAMPLING STEPS

We first study how the number of diffusion sampling steps affects ECG generation quality when keeping the sampler (DDPM) and noise schedule (linear) fixed. As shown in Table 6, performance generally improves as we increase the number of steps: reducing the steps from 1000 to 800 and 500 leads to progressively higher NRMSE, larger heart-rate error ( $\text{MAE}_{\text{HR}}$ ), and lower rCLIP scores.

1026 The default setting of 1000 steps achieves the best trade-off among these metrics, and is therefore  
 1027 used in our main experiments; smaller step counts can still be used in latency-constrained scenarios  
 1028 at the cost of degraded fidelity and semantic alignment.

1030 Table 6: ECG generation performance under different sampling steps.  
 1031

Sampling Steps	NRMSE $\downarrow$	MAE <sub>HR</sub> $\downarrow$	rCLIP Score $\uparrow$
1000	<b>0.0714 <math>\pm</math> 0.0010</b>	<b>8.43 <math>\pm</math> 0.42</b>	<b>0.9470 <math>\pm</math> 0.0029</b>
800	0.0877 $\pm$ 0.0011	13.59 $\pm$ 1.05	0.8671 $\pm$ 0.0040
500	0.1097 $\pm$ 0.0010	15.60 $\pm$ 1.19	0.7888 $\pm$ 0.0084

1032 Next, we compare different sampler types while fixing the number of sampling steps (1000) and  
 1033 using a linear noise schedule. Table 7 shows that our default DDPM sampler in SE-Diff achieves  
 1034 the best overall balance across MAE<sub>HR</sub>, and rCLIP. While the SDE-based sampler attains slightly  
 1035 lower NRMSE, it substantially worsens heart-rate accuracy and semantic alignment (higher MAE<sub>HR</sub>  
 1036 and lower rCLIP). DDIM and the second-order solver also underperform DDPM in at least one key  
 1037 metric. These results support our choice of DDPM as the default sampler for ECG generation in  
 1038 SE-Diff.

1044 Table 7: ECG generation performance under different samplers.  
 1045

Sampler	NRMSE $\downarrow$	MAE <sub>HR</sub> $\downarrow$	rCLIP Score $\uparrow$
<b>SE-Diff (ours)</b>	0.0714 $\pm$ 0.0010	<b>8.43 <math>\pm</math> 0.42</b>	<b>0.9470 <math>\pm</math> 0.0029</b>
DDIM	0.1219 $\pm$ 0.0010	9.05 $\pm$ 0.44	0.7783 $\pm$ 0.0072
Second Order	0.0877 $\pm$ 0.0011	13.59 $\pm$ 1.05	0.8671 $\pm$ 0.0040
SDE	<b>0.0693 <math>\pm</math> 0.0106</b>	19.21 $\pm$ 1.27	0.9213 $\pm$ 0.0028

1054 Finally, we investigate the impact of the noise schedule under the same DDPM sampler and 1000  
 1055 sampling steps. As reported in Table 8, both the linear and cosine schedules yield comparable  
 1056 performance, but the linear schedule achieves slightly better NRMSE, MAE<sub>HR</sub>, and rCLIP. This  
 1057 confirms that our default linear schedule is a strong and stable choice for ECG generation in SE-  
 1058 Diff, while alternative schedules can be used if desired without dramatically changing performance.

1059 Table 8: ECG generation performance under different noise schedules.  
 1060

Sampler	NRMSE $\downarrow$	MAE <sub>HR</sub> $\downarrow$	rCLIP Score $\uparrow$
Linear (ours)	<b>0.0714 <math>\pm</math> 0.0010</b>	<b>8.43 <math>\pm</math> 0.42</b>	<b>0.9470 <math>\pm</math> 0.0029</b>
Cosine	0.0728 $\pm$ 0.0032	8.90 $\pm$ 0.78	0.9446 $\pm$ 0.0131

1069 

## J ROBUSTNESS OF RCLIP AND RFID TO ENCODER CHOICES

  
 1070

1071 To assess the potential dependence of rCLIP and rFID on the particular ECG/text encoder pair, we  
 1072 compare SE-Diff and DiffuSETS under three encoder configurations. **Config A** uses our baseline  
 1073 ECG and text encoders. **Config B** replaces the ECG encoder with an alternative ECG representation  
 1074 while keeping the text encoder fixed. **Config C** instead replaces the text encoder while keeping the  
 1075 ECG encoder fixed. For each configuration, we jointly report rCLIP, rFID, signal-level reconstruc-  
 1076 tion error (MAE), and heart-rate error (MAE<sub>HR</sub>).

1077 As shown in Table 9, SE-Diff consistently outperforms DiffuSETS across *all* encoder configurations  
 1078 and metrics: MAE and MAE<sub>HR</sub> are always lower for SE-Diff, and both rCLIP and rFID are always  
 1079 higher. Moreover, the relative ranking between SE-Diff and DiffuSETS is stable when switching  
 ECG or text encoders, indicating that our conclusions are not tied to a specific encoder choice. We

1080 also observe that improvements in rCLIP/rFID co-occur with gains in heart-rate and signal-level  
 1081 fidelity, suggesting that these representation-based metrics are aligned with clinically meaningful  
 1082 waveform quality rather than being purely encoder-specific artifacts.  
 1083

1084 **Table 9:** ECG generation performance comparison between SE-Diff and DiffuSETS under different  
 1085 calibration configurations.

1087 Config	1088 Model	1089 MAE $\downarrow$	1090 MAE <sub>HR</sub> $\downarrow$	1091 rCLIP Score $\uparrow$	1092 rFID Score $\uparrow$
1088 A	1089 DiffuSETS	1090 $0.1092 \pm 0.0022$	1091 $13.29 \pm 1.13$	1092 $0.9309 \pm 0.0036$	1093 $0.9209 \pm 0.0038$
1089 A	1090 <b>SE-Diff (ours)</b>	1091 $0.0923 \pm 0.0021$	1092 $8.43 \pm 0.42$	1093 $0.9470 \pm 0.0029$	1094 $0.9509 \pm 0.0035$
1090 B	1091 DiffuSETS	1092 $0.1092 \pm 0.0022$	1093 $13.29 \pm 1.13$	1094 $0.9347 \pm 0.0032$	1095 $0.9197 \pm 0.0037$
1091 B	1092 <b>SE-Diff (ours)</b>	1093 $0.0923 \pm 0.0021$	1094 $8.43 \pm 0.42$	1095 $0.9472 \pm 0.0032$	1096 $0.9505 \pm 0.0035$
1092 C	1093 DiffuSETS	1094 $0.1092 \pm 0.0022$	1095 $13.29 \pm 1.13$	1096 $0.9242 \pm 0.0039$	1097 $0.9209 \pm 0.0038$
1093 C	1094 <b>SE-Diff (ours)</b>	1095 $0.0923 \pm 0.0021$	1096 $8.43 \pm 0.42$	1097 $0.9393 \pm 0.0031$	1098 $0.9509 \pm 0.0035$

## 1097 K SENSITIVITY ANALYSIS OF EXPERIENCE RETRIEVAL–AUGMENTED 1098 CONDITIONING

1100 Table 10 summarizes the sensitivity of the retrieval performance to the tri-view weights ( $\lambda_1, \lambda_2, \lambda_3$ )  
 1101 (diagnoses / prescriptions / procedures) and the number of retrieved admissions  $k$ . With equal  
 1102 weights ( $\lambda_1, \lambda_2, \lambda_3$ ) = (1, 1, 1), varying  $k \in 2, 3, 4$  yields very similar mean scores (94.4–94.8),  
 1103 with a slight peak at  $k = 3$  ( $94.9 \pm 2.3$ ), indicating limited sensitivity to  $k$  in this range. For  $k = 3$ ,  
 1104 varying one weight at a time while keeping the others at 1 shows that upweighting any single view  
 1105 changes performance by at most  $\sim 1.5$  points, suggesting that the tri-view retrieval is reasonably  
 1106 robust and that the symmetric setting ( $\lambda_1, \lambda_2, \lambda_3$ ) = (1, 1, 1) with  $k = 3$  is a simple and competitive  
 1107 default used in all main experiments. Retrieval performance is quantified by asking an LLM (GPT-  
 1108 5) to score each retrieved experience report on a 1–100 scale and reporting the mean  $\pm$  standard  
 1109 deviation of these scores.

1110 **Table 10:** Sensitivity of tri-view weights ( $\lambda_1, \lambda_2, \lambda_3$ ) and top- $k$  on retrieval performance (Mean  $\pm$   
 1111 Std).

1113 $\lambda_1$	1114 $\lambda_2$	1115 $\lambda_3$	1116 $k$	1117 <b>Mean <math>\pm</math> Std</b>
1115 1	1116 1	1117 1	1118 3	1119 <b><math>94.9 \pm 2.3</math></b>
1116 1	1117 1	1118 2	1119 2	1120 $94.4 \pm 2.1$
1117 1	1118 1	1119 4	1120 4	1121 $94.8 \pm 2.9$
1118 2	1119 1	1120 1	1121 3	1122 $94.7 \pm 2.6$
1119 1	1120 2	1121 1	1122 3	1123 $93.6 \pm 1.5$
1120 1	1121 1	1122 2	1123 3	1124 $93.5 \pm 3.3$
1121 3	1122 1	1123 1	1124 3	1125 $94.3 \pm 2.0$
1122 1	1123 3	1124 1	1125 3	1126 $94.6 \pm 1.6$
1123 1	1124 1	1125 3	1126 3	1127 $94.2 \pm 1.4$

## 1127 L ILLUSTRATIVE EXAMPLE OF EXPERIENCE RETRIEVAL–AUGMENTED 1128 CONDITIONING

1130 To make the experience retrieval–augmented conditioning pipeline more concrete, we provide a  
 1131 representative example below.

- 1132 • **Original machine-generated diagnosis and metadata.**  
 1133 Original machine-generated diagnosis: *Sinus rhythm*.

1134 Heart rate (bpm): 83.  
1135 Age (years): 91.  
1136 Gender: F.

- 1137 • **Top- $k$  retrieved similar admissions (diagnosis view,  $k = 3$ ).**

1138     *Sinus rhythm*

1139     *Sinus rhythm*

1140     *Sinus rhythm with borderline 1st degree A–V block*

- 1141 • **LLM-generated experience report.**

1142     The ECG demonstrates a sinus rhythm at a rate of 83 beats per minute in a 91-  
1143     year-old female. The rhythm is regular with normal P wave morphology preced-  
1144     ing each QRS complex, consistent with sinus node origin. There is no evidence  
1145     of atrioventricular conduction delay or block; PR interval is within normal lim-  
1146     its. QRS complexes are narrow, indicating normal ventricular conduction. No  
1147     arrhythmias or ectopic beats are noted. Overall, the tracing reflects a stable si-  
1148     nus rhythm appropriate for the patient’s age without signs of sinus tachycardia or  
1149     conduction abnormalities.

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187